user guide

hp StorageWorks modular SAN array 1000

Fifth Edition (April 2003)

Part Number: 230941-005

The HP StorageWorks Modular SAN Array 1000 is a 2-Gb Fibre Channel storage system designed for the entry-level to mid-range Storage Area Network (SAN). Designed to reduce the complexity, expense, and risk of SAN deployments in heterogeneous configurations, the MSA1000 provides a scalable, high-performance system.

© 2001-2003 Hewlett-Packard Development Company, L.P.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice.

Microsoft®, Windows®, and Windows NT® are registered trademarks of Microsoft Corporation.

Intel® and Celeron® are registered trademarks of Intel Corporation.

UNIX® is a registered trademark of The Open Group.

Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Hewlett-Packard Company products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

Modular SAN Array 1000 User Guide
Fifth Edition (April 2003)
230941-005

contents

Overview xiv
Intended Audience xiv
Related Documentation xiv
Conventions xv
Document Conventions xv
Text Symbols xvi
Equipment Symbols xvi
Rack Stability xviii
Getting Help xviii
HP Technical Support xviii
HP Website xix
HP Authorized Reseller xix

1 Introduction

About the MSA1000 2
Features of the MSA1000 3
Supported Software Components 4
Front View of the MSA1000 5
Rear View of the MSA1000 6
Hot-Pluggability 7
System Power 8
Maximum Storage 8
SCSI ID Assignments 8
MSA1000 Hardware Components 9
MSA1000 Controller 9
MSA1000 Controller Display 10
Array Accelerator (Battery-backed Cache) 11
Array Configuration 13
SCSI I/O Module with an Integrated Environmental Monitoring Unit (EMU) 15
Redundant Power Supply/Blower Assembly 16

Fibre Channel I/O Module 17
MSA SAN Switch 2/8 18
MSA Hub 2/3 19
MSA1000 Software Components 20
Array Configuration Utility (ACU) 20
Command Line Interface 21
Insight Manager 7 21
Compaq Analyze 22

2 Operating System Specific Information
Summary of Supported Software Components 24
Microsoft Windows Environments 25
Supported Operating System Versions 25
Supported Software Components 25
Supported Host Bus Adapters 26
Specific Configuration Needs 26
Linux Environments 26
Supported Operating System Versions 26
Supported Software Components 26
Supported Host Bus Adapters 27
Specific Configuration Needs 27
Novell NetWare Environments 27
Supported Operating Systems 27
Supported Software Components 28
Supported Host Bus Adapters 28
Specific Configuration Needs 28
OpenVMS Environments 28
Supported Operating System Versions 28
Supported Software Components 29
Supported Host Bus Adapters 29
Specific Configuration Needs 29
Tru64 UNIX Environments 31
Supported Operating System Versions 31
Supported Software Components 31
Supported Host Bus Adapters 31
Specific Configuration Needs 31

3 MSA1000 Setup and Sample Configurations
MSA1000 Setup 34

Installing and Connecting Hardware Components and Cables 34
Installing MSA1000 Option Kits. 35
Installing the MSA1000 in the Rack 35
Installing the Host Bus Adapters in the Servers 36
Connecting the Cables. 36
Installing Support Software 39
Configuring the Switches, Routers, and Servers 40
Setting up and Configuring the Storage. 40
Determining a Storage Plan. 40
Configuring the Storage Arrays and Logical Volumes 41
Setting up and Configuring Secure Path in Redundant Configurations 41
Identifying the Server Host Mode 42
MSA1000 Sample Configurations 42
Fibre Channel I/O Module Configurations 43
Embedded MSA SAN Switch 2/8 Configurations. 46
Embedded MSA Hub 2/3 Configurations 48

4 Operation and Management
Connecting the Power. 50
Applying Power 51
Interpreting Component Indicators. 52
MSA1000 Controller Indicators 52
Enclosure Status Indicators 54
Power Supply/Blower Assembly Indicators 55
SCSI I/O Module with Integrated EMU Indicators 56
Hard Drive Indicators. 57
Fibre Channel I/O Module Indicators 59
Replacing Components. 61
Replacing the MSA1000 Controller 62
Replacing the MSA1000 Controller Cache 64
Replacing the Controller Cache Battery Pack 67
Replacing a Variable Speed Blower 71
Replacing a Power Supply 73
Replacing Hard Drives. 74
Replacing the 2-Gb Small Form Factor Pluggable (SFP) Transceiver 78
Adding Additional Storage Enclosures 79
Migrating Existing Storage Enclosures 79
Adding a New Storage Enclosure 81
Updating the Controller Firmware 82

5 Array Configuration Utility (ACU)

Installing the ACU on the Server 84
Accessing the ACU. 85
Using the Smart Component 85
Using Insight Manager 7 87
Description of Screen Regions 88
Typical Wizard-Based Screen 88
Typical Standard Mode Screen 89
Configuring a New Controller 90
Using the Express Configuration Mode to Configure a New Controller. 91
Using the Configuration Wizards to Configure a New Controller 93
Creating an Array 93
Creating a Logical Drive 96
Using the Standard Configuration Mode to Configure a New Controller 98
Modifying an Existing Controller. 100
Using the Express Configuration Mode to Modify an Existing Controller 100
Using the Configuration Wizards to Modify an Existing Controller. 101
Clear Configuration 101
Controller Settings. 102
Create an Array 103
Create a Logical Drive 104
Delete Arrays 105
Delete Logical Drives 106
Expand Array 106
Extend Logical Drive 108
Migrate a Logical Drive 109
Spare Management 110
Selective Storage Presentation (SSP) 111
Using Standard Configuration Mode to Modify an Existing Controller 112
Probability of Logical Drive Failure. 114

6 Command Line Interface (CLI)

CLI Overview 118
CLI Command Syntax 119
Overview of CLI Storage Configuration Procedures. 120
CLI Setup 121
Setting up the Serial Connection 121
Help Commands 124
Displaying a List of All Basic Commands. 124

Displaying a List of All Available Commands . . . 125
Displaying Help for a Specific Command . . . 126
Display Commands . . . 127
Viewing Physical Disk Information . . . 127
Viewing LUN Information . . . 128
Viewing LUN Names . . . 130
Viewing Controller Settings . . . 130
Viewing Global Settings . . . 132
Viewing Version Information . . . 133
Array Controller Configuration Commands . . . 134
Setting Global Variables . . . 134
Setting the Controller ID . . . 135
Changing the CLI Prompt . . . 136
LUN Management Commands . . . 137
Flashing LEDs/Locating Hard Drives . . . 137
Locating all Drives Attached to the MSA1000 . . . 138
Locating all Drives in a Specific Storage Enclosure . . . 138
Locating all Drives for a Specific SCSI Bus . . . 139
Locating all Drives in a Specific LUN . . . 139
Locating Specific Drives . . . 139
Canceling a LOCATE Request . . . 139
Creating LUNs . . . 140
Creating a Single LUN from a Group of Drives . . . 141
Creating Multiple LUNs on a Group of Drives . . . 142
Creating a LUN with an Assigned Spare . . . 144
Assigning a Name or ID to a LUN . . . 145
Adding a Spare to a LUN . . . 145
Deleting LUNs . . . 146
Deleting Spares . . . 147
Recognizing a Failed Unit . . . 147
Modifying Arrays and LUNs . . . 148
Adding Drives to an Array . . . 148
Adding Space to a LUN . . . 149
Changing the RAID Characteristics of a LUN . . . 150
Changing Attributes of a LUN . . . 151
Server Connection Commands . . . 152
Viewing Existing Connections and the HBA World Wide Name . . . 153
Naming a Connection to a server . . . 154

Managing Connection Profiles. 155
Selective Storage Presentation/Access Control List Commands 156
Viewing the Connection Names. 157
Viewing the ACL. 157
Adding to the ACL. 157
Adding to the ACL using the Connection Name. 158
Adding to the ACL using the WWPN. 158
Deleting Information from the ACL 159
Deleting Information from the ACL using the connection name 159
Deleting Information from the ACL using the WWPN 160
Changing the Name of a Connection. 160
Changing the HBA of a Connection 161
Deleting a Connection Name. 161
Disabling the ACL. 162

A Regulatory Compliance Notices

Regulatory Compliance Identification Numbers 163
Federal Communications Commission Notice. 163
Modifications. 163
Cables 163
Canadian Notice (Avis Canadien) 164
European Union Notice 164
Japanese Notice 164
BSMI Notice 165
Laser Compliance. 165
Battery Replacement Notice 166

B Electrostatic Discharge

Grounding Methods 168

C Specifications

D Hard Drive Arrays

What is a Drive Array?. 172
Fault-Tolerance Methods 176
RAID 0—No Fault Tolerance 176
Advantages 176
Disadvantages 176
RAID 1 Drive Mirroring/RAID 1+0—Data Striping and Drive Mirroring. 177

Advantages 178
Disadvantages 178
RAID 5—Distributed Data Guarding 179
Advantages 179
Disadvantages 179
RAID ADG—Advanced Data Guarding 180
Advantages 180
Disadvantage 180
Other Fault-Tolerance Options 182

E Recovering from Hard Drive Failure
Hard Drive Failure 184
Recognizing Drive Failure 185
Compromised Fault Tolerance 186
Procedure to Attempt Recovery 187
Automatic Data Recovery 188
Failure of Automatic Data Recovery 188
Replacing a Drive 189
Moving Array Controller Drives and Arrays 191
Expanding and Extending Capacity 192

F Controller Display Messages
About the LCD Messages 196
Error Messages 196
Informational Messages 196
User Input Messages 196
Scrolling 197
User Input 197
Deleting Messages 197
Redundancy Link Light 197
LCD Message Descriptions 198

G Recovery ROM and ROM Cloning
Recovery ROM 229
ROM Cloning 230

H SCSI ID Assignments
SCSI ID Assignments for the MSA1000 232
SCSI ID Assignments for 4314 Storage Enclosures 233

SCSI ID Assignments for 4354 Storage Enclosures. 234

Figures

1 Front view of the MSA1000 5
2 Rear view of the MSA1000 6
3 Controller display 10
4 Basic non-cooperating server configuration using the switch. 14
5 SCSI I/O Module with an integrated Environmental Monitoring Unit 15
6 Power supply/blower assembly 16
7 Fibre Channel I/O Module. 17
8 MSA SAN Switch 2/8 18
9 MSA Hub 2/3. 19
10 2 Gb/s and 1 Gb/s connectors 37
11 Cable management using cable ties 1 38
12 I/O module directly connected to a server. 43
13 I/O module connected to an external switch, non-redundant configuration 44
14 I/O module connected to an external SAN switch, redundant configuration 45
15 MSA SAN Switch 2/8, non-redundant configuration 46
16 MSA SAN Switch 2/8, redundant configuration. 47
17 MSA Hub 2/3 connected to two servers, redundant configuration 48
18 MSA1000 Controller indicators 52
19 Enclosure status indicators 54
20 Power supply/blower assembly indicators 55
21 SCSI I/O Module with Integrated EMU 56
22 Hard drive indicators 57
23 Fibre Channel I/O module indicators 59
24 Removing the MSA1000 Controller 62
25 Installing the replacement controller 63
26 Removing the controller 64
27 Removing the cache module 65
28 Installing the cache module 66
29 Installing the controller 66
30 Bottom clip on battery pack. 68
31 Angling the battery pack 68
32 Removing the battery pack 69
33 Installing the battery pack 70
34 Securing the battery pack 70
35 Replacing the blower 71

36 Installing the replacement blower 72
37 Removing the power supply 73
38 Installing the replacement power supply 74
39 Opening the drive tray latches 76
40 Inserting the new drive 77
41 Removing the failed SFP 78
42 MSA1000 expansion SCSI connectors 80
43 ProLiant Storage Manager welcome screen 85
44 Configuration mode selection screen 86
45 Typical wizard-based screen 88
46 Typical More Information popup window 89
47 Typical standard mode screen 89
48 Starting to configure a new controller 90
49 Express mode start screen 91
50 Choosing a RAID level 92
51 Physical drive selection screen 93
52 Configured array screen (no logical drives) 95
53 New logical drive before saving 97
54 Typical task list for a logical drive 98
55 Typical SSP screen for MSA 1000 111
56 Typical task list for a logical drive 112
57 Relative probability of a logical drive failure 115
58 HyperTerminal New Connection dialog box. 122
59 HyperTerminal Connect To dialog box 122
60 HyperTerminal COM1 Properties dialog box 123
61 Physical drives added to system 172
62 Physical drives configured into a logical drive (L1) 173
63 Data striping (S1-S4) of data blocks B1-B12 173
64 Two arrays (A1, A2) containing five logical drives 174
65 Drive mirroring of P1 to P2 177
66 Data Striping and Drive Mirroring 178
67 Distributed data guarding, showing parity information (Px,y) 179
68 Advanced data guarding (RAID ADG) 180
69 Hard drive indicators 185
70 Array capacity expansion 193

Tables

1 Document Conventions xv

2	MSA1000 Features, summarized	3
3	MSA1000 Supported Software	4
4	Available Array Configuration Methods by Operating System	21
5	Supported Software Components, Summarized	24
6	Supported HBAs	36
7	Configuration Tools	41
8	MSA1000 Controller Indicator Descriptions	53
9	Determining Hard Drive Status from Indicator Combinations	57
10	Fibre Channel I/O Module Global Status Indicator	59
11	Fibre Channel I/O Module Link Status Indicators	60
12	Optimum Stripe Size	96
13	Possible Tasks in Standard Configuration Mode	99
14	Optimum Stripe Size	104
15	Possible Tasks in Standard Configuration Mode	113
16	CLI Special Keys	118
17	MSA1000 Specifications	169
18	Summary of RAID Methods	181
19	Choosing a RAID Method	182
20	LCD Message Descriptions	198
21	MSA1000 SCSI ID Assignments	232
22	Additional 4314 Storage Enclosure SCSI ID Assignments	233
23	Additional 4354 Storage Enclosure	234

about this guide

This user guide provides information to help you:

- Operate the MSA1000
- Configure the MSA1000

About this Guide topics include:

- Overview, page xiv
- Conventions, page xv
- Rack Stability, page xviii
- Getting Help, page xviii

Overview

This section covers the following topics:

- Intended Audience
- Related Documentation

Intended Audience

This book is intended for use by administrators with a moderate amount of SAN-management experience.

Related Documentation

In addition to this guide, refer to the HP StorageWorks Modular SAN Array 1000 Installation Overview poster that ships with this system.

Conventions

Conventions consist of the following:

- Document Conventions
- Text Symbols
- Equipment Symbols

Document Conventions

The document conventions included in Table 1 apply in most cases.

Table 1: Document Conventions

Element	Convention
Cross-reference links	Figure 1
Key and field names, menu items, buttons, and dialog box titles	**Bold**
File names, application names, and text emphasis	*Italics*
User input, command and directory names, and system responses (output and messages)	`Monospace font` `COMMAND NAMES` are uppercase monospace font unless they are case sensitive
Variables	`<monospace, italic font>`
Website addresses	Underlined sans serif font text: http://www.hp.com

Text Symbols

The following symbols may be found in the text of this guide. They have the following meanings.

WARNING: Text set off in this manner indicates that failure to follow directions in the warning could result in bodily harm or loss of life.

Caution: Text set off in this manner indicates that failure to follow directions could result in damage to equipment or data.

Note: Text set off in this manner presents commentary, sidelights, or interesting points of information.

Equipment Symbols

The following equipment symbols may be found on hardware for which this guide pertains. They have the following meanings.

Any enclosed surface or area of the equipment marked with these symbols indicates the presence of electrical shock hazards. Enclosed area contains no operator serviceable parts.

WARNING: To reduce the risk of injury from electrical shock hazards, do not open this enclosure.

Any RJ-45 receptacle marked with these symbols indicates a network interface connection.

WARNING: To reduce the risk of electrical shock, fire, or damage to the equipment, do not plug telephone or telecommunications connectors into this receptacle.

Any surface or area of the equipment marked with these symbols indicates the presence of a hot surface or hot component. Contact with this surface could result in injury.

WARNING: To reduce the risk of injury from a hot component, allow the surface to cool before touching.

Power supplies or systems marked with these symbols indicate the presence of multiple sources of power.

WARNING: To reduce the risk of injury from electrical shock, remove all power cords to completely disconnect power from the power supplies and systems.

Any product or assembly marked with these symbols indicates that the component exceeds the recommended weight for one individual to handle safely.

WARNING: To reduce the risk of personal injury or damage to the equipment, observe local occupational health and safety requirements and guidelines for manually handling material.

Rack Stability

Rack stability protects personnel and equipment.

WARNING: To reduce the risk of personal injury or damage to the equipment, be sure that:

- The leveling jacks are extended to the floor.
- The full weight of the rack rests on the leveling jacks.
- In single rack installations, the stabilizing feet are attached to the rack.
- In multiple rack installations, the racks are coupled.
- Only one rack component is extended at any time. A rack may become unstable if more than one rack component is extended for any reason.

Getting Help

If you still have a question after reading this guide, contact an HP authorized service provider or access our website: http://www.hp.com.

HP Technical Support

In North America, call technical support at 1-800-652-6672, available 24 hours a day, 7 days a week.

Note: For continuous quality improvement, calls may be recorded or monitored.

Outside North America, call technical support at the nearest location. Telephone numbers for worldwide technical support are listed on the HP website under support: http://www.hp.com.

Be sure to have the following information available before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

HP Website

The HP website has the latest information on this product, as well as the latest drivers. Access storage at: http://www.hp.com. From this website, select the appropriate product or solution.

HP Authorized Reseller

For the name of your nearest HP Authorized Reseller:

- In the United States, call 1-800-345-1518
- In Canada, call 1-800-263-5868
- Elsewhere, see the HP website for locations and telephone numbers: http://www.hp.com.

Introduction

1

The HP StorageWorks Modular SAN Array 1000 (MSA1000) is a next generation, 2-Gb storage system designed for the entry level to mid-range Storage Area Network (SAN). Designed to reduce the complexity, expense, and risk of SAN deployment, the MSA1000 provides a scalable, high-performance storage system with investment protection in mind. The modular design of the MSA1000 allows storage capacity to be added as needed.

This introductory chapter includes the following overviews:

- About the MSA1000, page 2
- MSA1000 Hardware Components, page 9
- MSA1000 Software Components, page 20

About the MSA1000

The MSA1000 provides a high performance array controller in a 14-drive storage cabinet, capable of 2 Gb/s speeds over Fibre Channel. It may be expanded by adding two more storage cabinets for a total capacity of 42 Ultra 2 or Ultra3 SCSI drives. The MSA1000 is compatible with 1 Gb/s and 2 Gb/s hub and fabric switch interconnects and includes options for embedded switches and hubs.

Complete operating redundancy is supported, including redundancy while expanding disk and unit volumes.

The MSA1000 can operate in standalone, dual-node, or multi-node cluster environments, supporting multiple operating systems, including Microsoft Windows®, Novell NetWare, Linux®, OpenVMS®, or Tru64 UNIX®.

This section discusses:

- Features of the MSA1000
- Supported software components
- Front view of the MSA1000
- Rear view of the MSA1000
- Hot-pluggability
- System power
- Maximum storage
- SCSI ID assignments

Features of the MSA1000

Key features of the MSA1000 are summarized in Table 2.

Table 2: MSA1000 Features, summarized

Feature	Description
Form factors	HP Series rack, 4U design combining both the controller and disk shelf
Model	Rack-mountable
Hard drive types	HP Ultra2 and Ultra3 1" SCSI hard disk drives (maximum drive size: 146.8 GB)
Maximum number of drives	14 (1-inch); scalable to 42 drives with two optional storage expansion enclosures
Connection to host	2-Gb Fibre Channel connections, with support for 1-Gb and 2-Gb Fibre Channel infrastructures.
Operating systems	Windows Novell NetWare Linux OpenVMS Tru64 UNIX
Controller	Embedded MSA1000 Controller with battery-backed cache (256 MB; can be upgraded to 512 MB)
RAID configuration levels	RAID 0, RAID 1, RAID 1+0, RAID 5, and Advanced Data Guarding (ADG)
Option	Embedded 2 Gb/s MSA SAN Switch 2/8
Option	Embedded 2 Gb/s MSA Hub 2/3

Note: Refer to the MSA1000 website: www.hp.com/go/msa1000, for the most up-to-date information on supported components and operating system versions.

Supported Software Components

Supported software is summarized in Table 3.

Table 3: MSA1000 Supported Software

Software	Operating System
Array Configuration Utility (ACU)	Windows: online Linux: online NetWare: offline
Command Line Interface (CLI)	Windows Linux NetWare OpenVMS Tru64 UNIX
Microsoft Clustering Services (MSCS)	Windows
OpenView Virtual Replicator	Windows
SmartStart	Windows Linux NetWare
Compaq Analyze	OpenVMS Tru64 UNIX
Insight Manager 7	Windows Linux NetWare
HP StorageWorks Secure Path for Windows Workgroup Edition	Windows
HP StorageWorks Secure Path for NetWare	NetWare
HP StorageWorks Secure Path for Linux	Linux

Note: Refer to the MSA1000 website: www.hp.com/go/msa1000, for the most up-to-date information on supported software components

Front View of the MSA1000

The following figure and supporting table illustrate and define the front view features of the MSA1000.

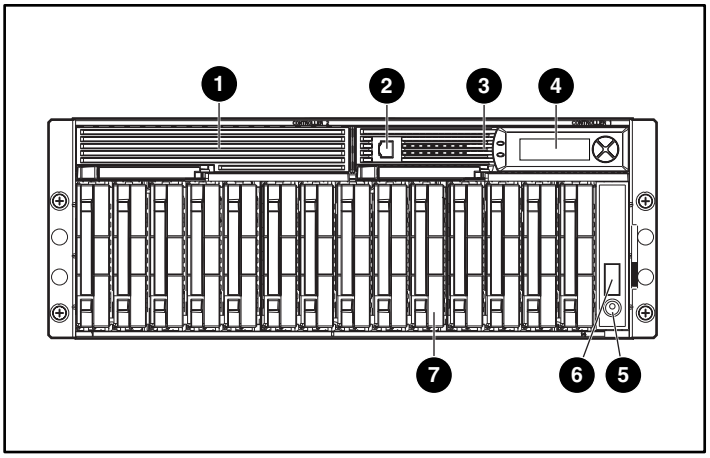

Figure 1: Front view of the MSA1000

Item	Description
❶	Blank panel Included in the standard shipping configuration and removed when a redundant MSA1000 Controller is installed. In redundant configurations, this is the secondary controller and may be referred to as controller2.
❷	RJ-45Z customized serial port for accessing the CLI One customized cable is included in the shipping carton.
❸	Hot-Pluggable MSA1000 Controller In redundant configurations, this is the primary controller and may be referred to as controller1.
❹	Controller LCD display
❺	Power switch
❻	Enclosure status indicators
❼	Hot-pluggable hard drives

Rear View of the MSA1000

The following figure and supporting table illustrate and define the features and components on the rear of the MSA1000.

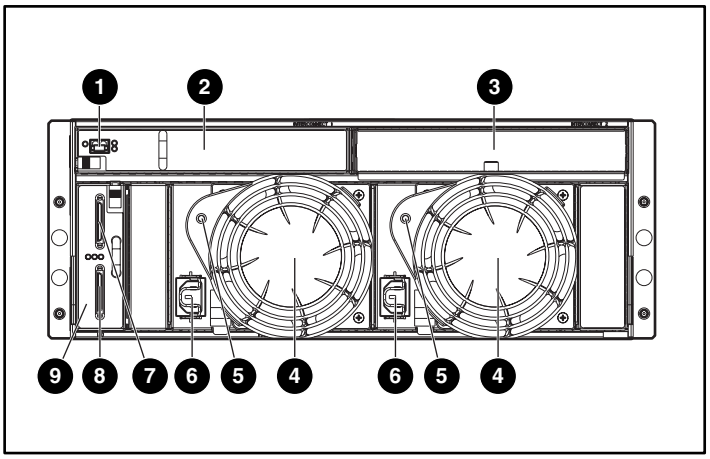

Figure 2: Rear view of the MSA1000

Item	Description
❶	2-Gb Small Form Factor Pluggable Transceivers (SFP)
❷	Fibre Channel I/O Module
❸	Blank panel (when hot-pluggable options are not installed)
❹	Power supply/blower assemblies
❺	Power supply/blower assembly indicators
❻	AC power cord connectors
❼	SCSI Port A connector
❽	SCSI Port B connector
❾	SCSI I/O Module with an integrated Environmental Monitoring Unit (EMU)

Hot-Pluggability

Hot-pluggability allows components to be removed and replaced without powering down the MSA1000.

The following MSA1000 hardware components are hot-pluggable:

Hard Drives—Failed drives in fault-tolerant configurations can be replaced without interrupting system operation. Data on the failed drive is restored automatically to the replacement drive online.

Power Supply/Blower Assembly—Redundant power supplies allow you to remove one power supply and replace it without interrupting system operation. The assembly includes one variable speed blower for each power supply. If one of the blowers in the assembly fails, the other blower automatically operates at a higher speed. A blower can be removed and replaced without removing the assembly.

MSA1000 Controller—If your system has two controllers and one controller fails, you can remove and replace the failed controller. See the "Operation and Management" chapter for specific limitations.

Fibre Channel I/O Module—The Fibre Channel I/O Module is a single-port, hot-pluggable 1-Gb/2-Gb Fibre Channel device that includes a 2-Gb Small Form Factor Pluggable Transceiver (SFP) port. This module can be removed and replaced with other interconnect options, such as the MSA SAN Switch 2/8 or the MSA Hub 2/3.

MSA SAN Switch 2/8—The MSA SAN Switch 2/8 is an optional, hot-pluggable I/O device designed to replace the single port Fibre Channel I/O Module that comes standard with the MSA1000. This switch includes 8-ports, featuring 2 Gb/s full duplex non-blocking performance on all ports. One port is internal to the switch, leaving 7 available external ports.

MSA Hub 2/3—The MSA Hub 2/3 is an optional, hot-pluggable I/O device designed to replace the single port Fibre Channel I/O Module that comes standard with the MSA1000. This hub provides an inexpensive and easy way for a clustered pair of servers to access the storage controller (using an arbitrated-loop configuration), without having to buy an external hub or switch. Two SFP ports are presented to the user, while a third, internal port accesses the controller.

System Power

Power to the MSA1000 does not shut off completely with the power switch. The two positions of the front panel power switch should be considered as *on* and *standby*, instead of *on* and *off*. The *standby* position removes power from most of the electronics and the drives, but portions of the power supply and some internal circuitry remain active.

Consequently, to remove all power from the MSA1000 storage system, you must disconnect the power cord from the storage system. In environments with multiple power supplies, you must disconnect all power cords to remove power completely from the MSA1000 storage system.

WARNING: To reduce the risk of electric shock or damage to the equipment, disconnect power from the storage system by unplugging all power cords from either the electrical outlet or the storage system.

Maximum Storage

A maximum of 14 1-inch Universal drives can be installed in the MSA1000. Up to two additional external storage enclosures can be cascaded from the MSA1000, expanding the capacity to 42 drives. Supported storage enclosures include the HP StorageWorks 4200 and HP StorageWorks 4300 product lines.

Supported drive sizes include 18 GB, 36.4 GB, 72.6 GB, and 146.8 GB. If the enclosures are fully populated with 146.8 GB drives, the capacity of the MSA1000 storage system is approximately 6 TB.

Note: The maximum size of a logical unit is 2 TB. If you have more than 2 TB of storage attached to your MSA1000, you must create multiple logical units.

SCSI ID Assignments

SCSI IDs are assigned automatically in the storage system according to the drive bay used for each drive. It is not necessary to assign SCSI IDs manually. Drive bay numbers are indicated on the front panel. See the "SCSI ID Assignments" appendix for a complete table of SCSI ID assignments.

MSA1000 Hardware Components

The standard configuration of the MSA1000 includes one MSA1000 Controller, one MSA Fibre Channel I/O Module, one SCSI I/O with an integrated EMU, and redundant power supply/blower assemblies.

Option kits are available for redundant MSA1000 Controllers and Fibre Channel I/O Modules as well as for embedded switches and hubs.

Components of the MSA1000 include:

- MSA1000 Controller
- SCSI I/O with integrated Environmental Monitoring Unit
- Redundant power supply/blower assemblies
- Fibre Channel I/O Module
- MSA SAN Switch 2/8 (option only)
- MSA Hub 2/3 (option only)

MSA1000 Controller

The MSA1000 Controller is a drive array controller specifically designed for installation in the MSA1000. The MSA1000 comes equipped with one MSA1000 Controller installed. An additional controller for redundancy can be purchased separately.

To ensure uninterrupted service, two copies of the controller firmware are stored in Read Only Memory (ROM) on the controller. See the "Recovery ROM and ROM Cloning" appendix for more information.

Additional information about the following topics is included in this section:

- MSA1000 Controller Display
- Array Accelerator (Battery-backed Cache)
- Array Configuration

MSA1000 Controller Display

Each array controller in an MSA1000 contains an integrated Liquid Crystal Display (LCD). This module is used for displaying informational and error messages, showing the status of the module, and for providing user input when required. Traditional Power-On Self-Test (POST) messages issued by PCI-based array controllers have been combined with runtime event notification messages to create a new set of controller display messages.

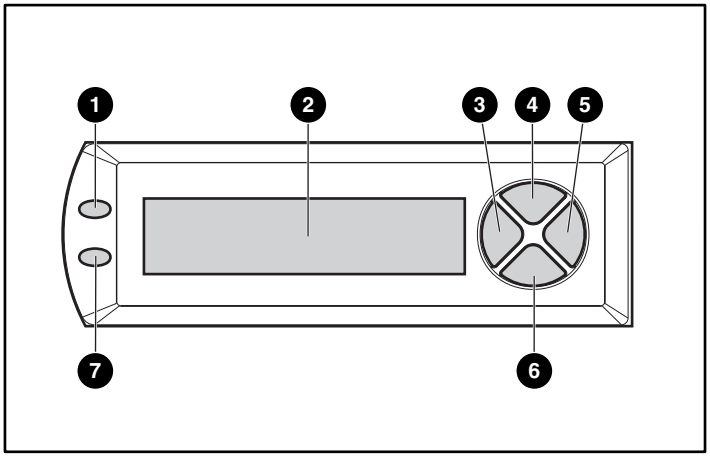

Figure 3: Controller display

Controller Display	Description
❶	Fault indicator (amber)
❷	Display
❸	Left push button
❹	Up push button
❺	Right push button
❻	Down push button
❼	Redundancy Link indicator (green)

For more information about the MSA1000 Controller display, see the "Controller Display Messages" appendix.

Array Accelerator (Battery-backed Cache)

The Array Accelerator is a high-performance, upgradeable 256-MB SDRAM DIMM read/write battery-backed cache that increases performance in database and fault-tolerant configurations. It performs both protected posted-write caching and read-ahead caching, allowing data to be accessed much faster than from disk storage.

In protected posted-write caching, data is written to the cache memory on the Array Accelerator rather than directly to the drives. Later, when the MSA1000 storage system is idle, the controller writes the cached data to the drive array.

The read-ahead cache detects sequential accesses to the array, reads ahead data, and stores the data in the cache until the next read access arrives. If the data is of a sequential nature, the data can be loaded immediately into memory, avoiding the latency of a disk access.

If the MSA1000 Controller fails before cached data is stored on the disk, the Array Accelerator and its integrated batteries may be removed from one MSA1000 Controller and installed on a replacement controller. Any data in the Array Accelerator that has not been written to the hard drive will be transferred to the replacement MSA1000 Controller.

Array Accelerator Features:

Other features of the Array Accelerator:

- Mounted on a removable daughterboard (allows stored data to be moved to another controller if the original controller fails)
- Backed up with replaceable batteries
- Upgradable to 512 MB (256 MB per controller)
- Adjustable read/write ratio - usually set during array configuration but can be changed at any time
- 16-bit Error Checking and Correcting (ECC) SDRAM memory

  ECC detects and corrects all single-bit memory errors. It also detects all two-bit memory errors in any position, and most three- and four-bit memory errors in a single SDRAM. With ECC, an entire memory chip can also fail without data loss. This provides a high level of data integrity by ensuring the correction of common memory errors without affecting performance.

Array Accelerator Batteries

The Array Accelerator has two rechargeable and replaceable NiMH battery packs. Under normal operating conditions, these should last for three years before replacement is necessary. They are continuously recharged via a "trickle" charging process whenever the MSA1000 storage system is powered on.

The battery packs protect data on the Array Accelerator against equipment failure or power outage for up to four continuous days.

Note: Temperature, age, and cache size may affect battery life.

This also applies if the Array Accelerator is removed from the MSA1000 Controller. When power is restored to the MSA1000 storage system, an initialization process writes the preserved data to the disk drives. This is particularly important for data that has been cached by a posted-write operation, but has not yet been written to the hard drives.

Note: The batteries on a new MSA1000 Controller may be discharged when the board is first installed. In this case, a Power-On Self-Test (POST) message will be displayed on the controller display panel when the controller is powered on, indicating that the Array Accelerator is temporarily disabled. No action is required on your part, since the internal circuitry will automatically recharge the batteries. Recharging the batteries can take up to 4 hours. The MSA1000 Controller will function properly during this time, although without the performance advantage of the Array Accelerator. When the batteries are charged to 90 percent of their capacity, the Array Accelerator is automatically enabled.

Depending on the status of the array accelerator, including a low battery charge, informational or error messages may be displayed on the controller's LCD panel. For a listing of cache module LCD messages, see the definitions for LCD messages 60 through 79 in the "Controller Display Messages" appendix.

For detailed array accelerator or battery replacement instructions, see the "Operation and Management" chapter.

Array Configuration

Drive array technology distributes data across a series of individual hard drives to unite these physical drives into one or more higher-performance logical arrays and volumes. Distributing the data allows for concurrent access from multiple drives in the array, yielding faster I/O rates than non-arrayed drives. Each logical drive in the array can be set to a different level of fault tolerance.

Depending on your server's operating system and your preference of user interfaces, use one of the following provided user interfaces to configure the storage on the MSA1000:

- Array Configuration Utility (ACU)
- Command Line Interface (CLI)

For information about these user interfaces, see the "MSA1000 Software Components" section also in this chapter.

Array Fault Tolerance Levels

Fault tolerance refers to several possible methods used to protect your MSA1000 storage system from data loss in the event of a hardware failure somewhere in the storage system. Each method has its advantages; choosing one requires careful planning for your particular requirements.

The RAID configuration methods supported by the MSA1000 Controller include:

- No fault tolerance (RAID 0)
- Drive mirroring (RAID 1)
- Drive mirroring and striping (RAID 1+0)
- Distributed data guarding (RAID 5)
- Advanced Data Guarding (RAID ADG)

Further data protection can be achieved by assigning one or more online spares to any fault-tolerant array.

For detailed information on fault tolerance options, see the "Hard Drive Arrays" appendix.

Selective Storage Presentation (SSP)

SSP permits sharing of the MSA1000 storage between servers, including servers running different operating systems.

Each logical drive on the controller has an Access Control List (ACL) that contains the worldwide names of the server host adapters that have access to the drive. If a server tries to send commands to a logical drive that it does not have access to, the firmware will reject the command.

The SSP maps the worldwide names of server host adapters to connection names and sets up ACLs for logical drives based on these worldwide names.

SSP can be set up and managed through the ACU and the CLI.

The following diagram illustrates an SSP scenario.

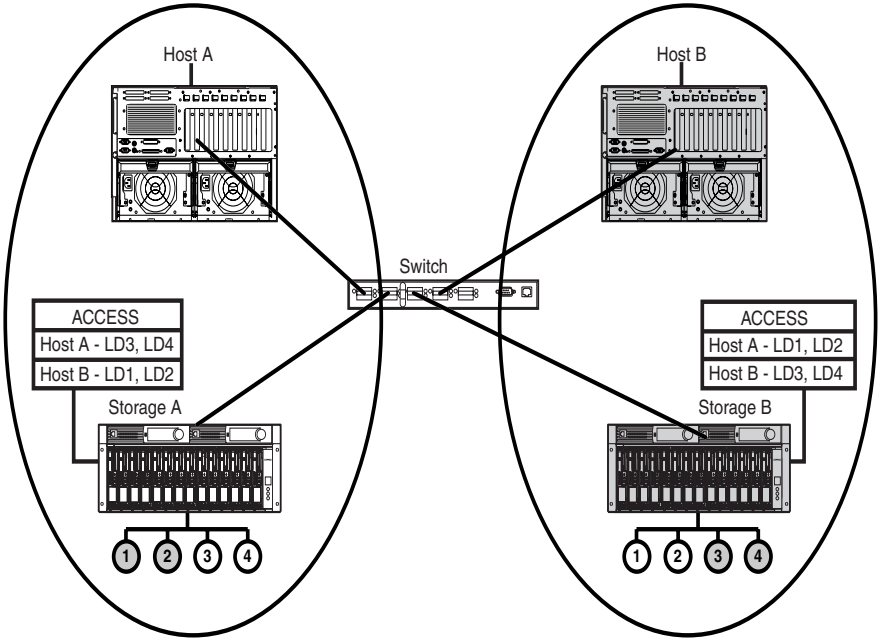

Figure 4: Basic non-cooperating server configuration using the switch

SCSI I/O Module with an Integrated Environmental Monitoring Unit (EMU)

This element performs several functions, including monitoring the enclosure operation, detecting and indicating errors, providing I/O connectors, and displaying enclosure status.

Figure 5 is an illustration of a SCSI I/O Module with the integrated EMU.

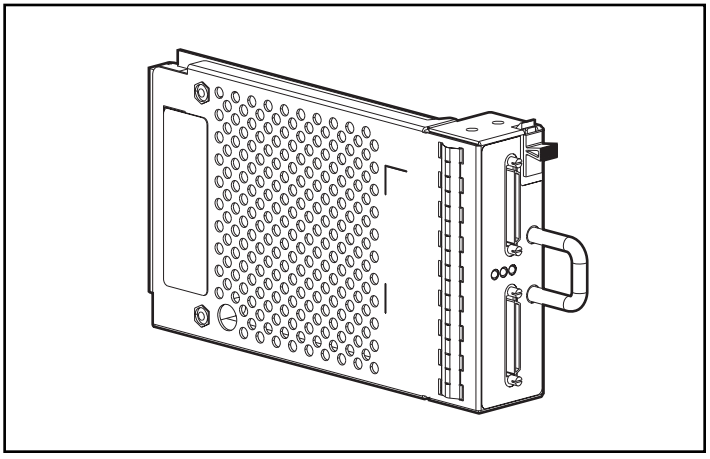

Figure 5: SCSI I/O Module with an integrated Environmental Monitoring Unit

EMU functions include:

- Monitoring the power supply voltages
- Monitoring the power supply and enclosure temperatures
- Monitoring the presence and status of hard drives
- Detecting and reporting element changes in the enclosure and showing the addition or removal of a hot-pluggable component
- Controlling drive status indicators
- Providing expansion ports for external HP StorageWorks enclosures

For more information on the EMU, including installation and removal and status indicators, refer to the "Enclosure Status Indicators" section of the "Operation and Management" chapter.

Redundant Power Supply/Blower Assembly

The MSA1000 includes a redundant power supply. The redundant power supply is identical to the primary supply and shares the electrical load equally. If either supply fails or is removed, the other power supply assumes the full load without interruption. Hot-pluggability allows you to replace the failed supply without shutting the MSA1000 storage system down.

The MSA1000 includes a redundant blower. If one blower operates too slowly, or completely shuts off, the other blower will begin to operate at a higher speed.

Any changes in the performance level of the power supply or blower are reported on the LCD panel of the associated controller. For a description of the possible informational and error messages associated with the power supplies or blowers, see the definitions for messages numbered 400 through 499 in the "Controller Display Messages" appendix.

Depending on the setup of Insight Manager 7 and Compaq Analyze, informational and error messages about the power supplies and blowers are also reported though these software components.

Figure 6 is an illustration of the power supply/blower assembly.

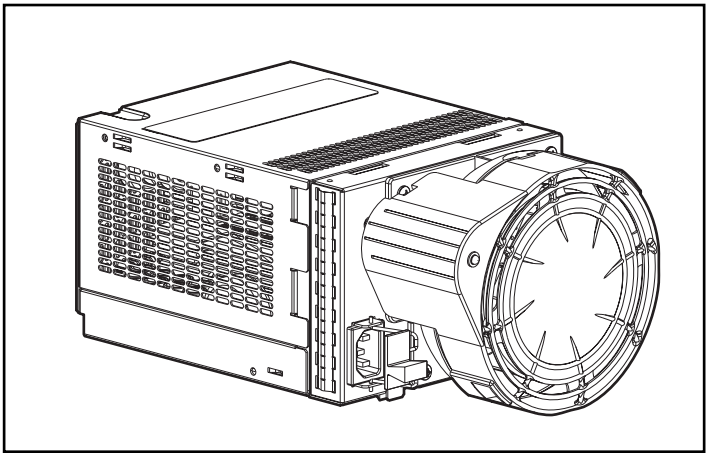

Figure 6: Power supply/blower assembly

Refer to the "Replacing a Power Supply" and "Replacing a Variable Speed Blower" sections in the "Operation and Management" chapter for more information.

Fibre Channel I/O Module

The Fibre Channel I/O Module is a single-port, 1-Gb/2-Gb Fibre Channel device that includes a 2 Gb/s Small Form Factor Pluggable Transceiver (SFP) port.

Figure 7 is an illustration of the Fibre Channel I/O Module.

This module can be removed and replaced with other interconnect options, such as the MSA SAN Switch 2/8 or the MSA Hub 2/3. In redundant configurations, the interconnect devices must be a matching pair.

If a second controller is added for additional functionality but not for system redundancy, an MSA SAN Switch 2/8 or an MSA Hub 2/3 may be installed.

If you are using the Fibre Channel I/O Module and an additional controller is installed for redundancy, a second I/O Module must be ordered and installed.

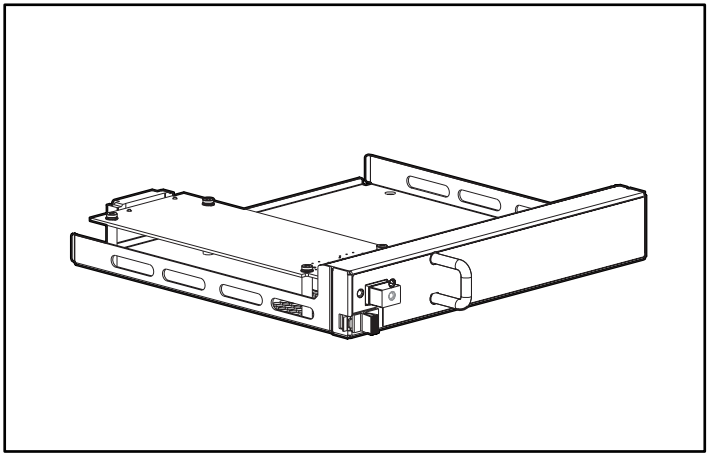

Figure 7: Fibre Channel I/O Module

MSA SAN Switch 2/8

The MSA SAN Switch 2/8 is an optional 2-Gb/s, integrated 8-port fabric switch (1 port internal to the MSA1000), that enables full duplex non-blocking performance on all ports. For information about the MSA SAN Switch 2/8, refer to the *HP StorageWorks MSA SAN Switch 2/8 Installation Guide*.

Figure 8 is an illustration of the MSA SAN Switch 2/8.

If you are using the MSA SAN Switch 2/8 and an additional controller is installed for redundancy, a second MSA SAN Switch 2/8 must be ordered and installed. In redundant configurations, the interconnect devices must be a matching pair.

If a second controller is added for additional functionality but not for system redundancy, an MSA SAN Switch 2/8 or an MSA Hub 2/3 may be installed.

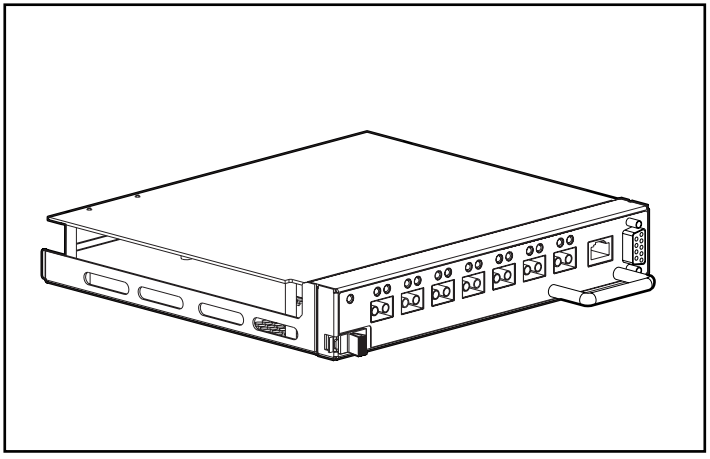

Figure 8: MSA SAN Switch 2/8

MSA Hub 2/3

The MSA Hub 2/3 is an optional input/output device designed to replace the single port Fibre Channel I/O Module that comes standard with the MSA1000. This option is designed to provide an inexpensive and easy way for a clustered pair of servers to access the storage controller (using the arbitrated loop configuration) without having to buy an external hub or switch. Two SFP ports are presented to the user while internally a third port accesses the controller.

Figure 9 is an illustration of the MSA Hub 2/3.

If a second controller is added for additional functionality but not for system redundancy, an MSA SAN Switch 2/8 or an MSA Hub 2/3 may be installed.

If you are using the MSA Hub 2/3 and an additional controller is installed for redundancy, a second MSA Hub 2/3 must be ordered and installed. In redundant configurations, the interconnect devices must be a matching pair.

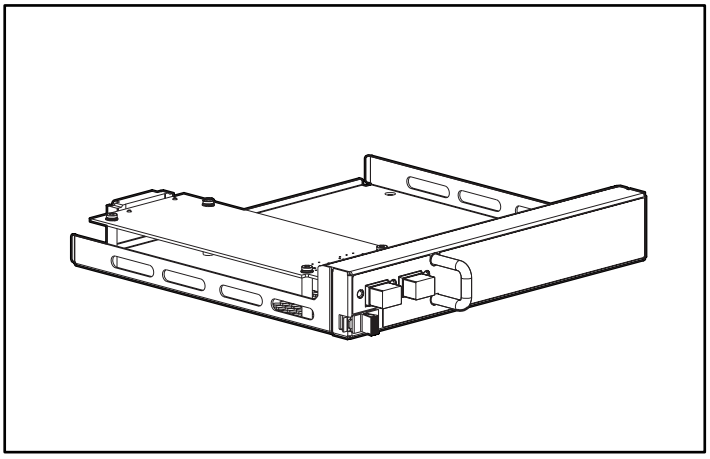

Figure 9: MSA Hub 2/3

MSA1000 Software Components

The following software components are supported on the MSA1000:

- Array Configuration Utility (ACU)
- Command Line Interface
- Insight Manager 7
- Compaq Analyze

Each software tool is discussed in the following paragraphs.

Array Configuration Utility (ACU)

The ACU is Web based and can run locally through the Windows Internet Explorer browser screen or remotely through Insight Manager 7. If the utility is not configured to allow remote access, it will run in the local environment.

The ACU is supported for Windows, Linux, and NetWare environments. In Windows and Linux environments, the ACU can be run online after it is installed on the server from the Modular SAN Array 1000 Support Software CD. NetWare environments must run the ACU from the CD.

All agent, network services, and management software must be installed, configured, and running before all of the features and functions of the ACU are operational.

For Web-browser or network access, a combination of the following may need to be installed and properly configured:

- Windows Internet Explorer
- TCP/IP and SNMP services
- Agents such as Insight Manger 7
- Host Bus Adapter drivers

Note: If you plan to use the ACU to configure and manage your storage, HP recommends using the ACU exclusively and not using the CLI.

For information on using the ACU, see the "Array Configuration Utility (ACU)" chapter.

Table 4: Available Array Configuration Methods by Operating System

Operating System	Array Configuration Method
Windows	ACU and CLI
Linux	ACU and CLI
NetWare	ACU (offline) and CLI
OpenVMS	CLI
Tru64 UNIX	CLI

Command Line Interface

The Command Line Interface (CLI) is built into the firmware and can be used to configure, manage, and monitor all aspects of the MSA1000, including hard drive array configuration. The CLI is accessed through a host computer connected to the customized RJ-45Z serial port on the front of the MSA1000.

All supported operating systems can access the CLI.

Note: If you plan to use the CLI to configure and manage your storage, HP recommends using the CLI exclusively and not using the ACU.

For information on using the CLI, see the "Command Line Interface (CLI)" chapter.

Insight Manager 7

Insight Manager 7 can be used to monitor the performance of several MSA1000 Controller parameters. The displayed parameters include CPU usage, the total number of read and write commands processed, and the average time to process a read or write command. Also displayed for each logical drive are the total I/O count, number of read and write requests, and number of sectors read or written. Use Insight Manager 7 to monitor the performance of your drive arrays.

The Windows, Linux, and NetWare operating systems can use Insight Manager 7.

For more information on Insight Manager 7, refer to the Web at http://www.hp.com.

Compaq Analyze

Compaq Analyze is a host-based fault analysis utility included with the Web-Based Service Enterprise (WEBES) service tools. Each WEBES service tool adds functionality to the Directory, a process (or set of processes) that executes continuously. Compaq Analyze provides the Director with the capability to capture and interpret hardware events. The analysis of events can be performed automatically or when requested by an outside process.

Compaq Analyze provides background automatic analysis by monitoring the active binary event log (also called the system event log or error log) and processing events as they occur. The events in the binary event log file are checked against the analysis rules. If one or more of the events in the binary event log file meets the conditions specified in the rules, the analysis engine collects the error data and creates a problem report containing a description of the problem and any corrective actions required. After the problem report is created, it is distributed in accordance with the customer's notification preferences.

Compaq Analyze features:

- Error event analysis and translation
- Multi-event correlation (view events, traps, and so on)
- Manual and automatic (real-time) analysis capabilities on local or remote machines
- Filtering capabilities on error logs
- Automatic notification (customer or HP Support Center)
- Rule-based (allowing more frequent updates)
- Proactive indictment with the operating system
- Web and command line interfaces

For more information, go to the HP website at http://www.hp.com and search on the keyword WEBES.

Operating System Specific Information

2

The MSA1000 supports several operating systems and can be deployed into heterogeneous environments with any combination of these operating systems.

Each operating system is introduced in a separate section of this chapter:

- Microsoft Windows Environments, page 25
- Linux Environments, page 26
- Novell NetWare Environments, page 27
- OpenVMS Environments, page 28
- Tru64 UNIX Environments, page 31

For each environment, information about the supported operating system, software components, and host bus adapters is detailed.

Summary of Supported Software Components

A variety of software components is supported on the MSA1000, with each operating system supporting a unique sub-set of these components.

Table 5 summarizes the availability of the supported software components by operating system.

As illustrated in Table 5, one primary difference between the operating systems is the choice of the storage configuration tool. The ACU or CLI may be available for a specific operating system.

Table 5: Supported Software Components, Summarized

Software Tool	Windows	Linux	NetWare	Open VMS	Tru64 UNIX
ACU	Y	Y	Y		
Compaq Analyze				Y	Y
Insight Manager 7	Y	Y	Y		
CLI	Y	Y	Y	Y	Y
Microsoft Cluster Service (MSCS)	Y				
HP StorageWorks Secure Path	Y	Y	Y	See note below	See note below
OpenView Virtual Replicator	Y				
SmartStart	Y	Y	Y		
NOTE: Tru64 UNIX and OpenVMS include equivalent capabilities in their base operating systems.					

Microsoft Windows Environments

Windows-based configurations that use the MSA1000 may range from a simple single-controller MSA1000 connected to a single node to a complex dual-node, redundant-controller MSA1000 connected to a complex fabric supporting multiple two-node clusters.

Supported Operating System Versions

- Windows NT 4.0 Server
- Windows 2000 Server
- Windows Server 2003

Note: For a current listing of the supported operating system versions, see the MSA1000 website at www.hp.com/go/msa1000.

Supported Software Components

- Array Configuration Utility
- Command Line Utility
- Insight Manager 7
- Microsoft Clustering Services (MSCS)
- HP StorageWorks Secure Path for Windows Workgroup Edition
- Storage Agents for Windows
- OpenView Virtual Replicator
- SmartStart and Support Software version 5.50 or later

Supported Host Bus Adapters

- FCA2101 Host Bus Adapter
- 64-bit/33-MHz PCI-to-Fibre Channel Host Bus Adapter

Specific Configuration Needs

Initial system configuration tasks are detailed in the Windows section of the *Modular SAN Array 1000 Quick Installation Guide*. This guide is included in the Modular SAN Array 1000 Setup and Management kit.

Linux Environments

Linux environments that use an MSA1000 can include multiple clusters.

When configuring the storage on the MSA1000, two methods are available: the CLI and the ACU are both available online.

Supported Operating System Versions

- Red Hat Linux 7.2 Professional
- Red Hat Linux Advanced Server 2.1
- SuSE Linux Enterprise Server 7

Note: For a current listing of the supported operating system versions, see the MSA1000 website at www.hp.com/go/msa1000.

Supported Software Components

- Array Configuration Utility
- Command Line Utility
- SteelEye clustering software
- Insight Manager 7
- HP StorageWorks Secure Path for Linux
- SmartStart version 5.50 or later

Supported Host Bus Adapters

- FCA2214 Host Bus Adapter

Specific Configuration Needs

Before connecting the MSA1000 to the server, install the FCA2214 HBA driver on the server to prepare the server for the MSA1000.

For instructions, see the Linux section in the *Modular SAN Array 1000 Quick Installation Guide*. This guide is included in the Modular SAN Array 1000 Setup and Management kit.

Novell NetWare Environments

Novell NetWare supports up to six nodes in a cluster using single paths. Complete redundancy with multiple clusters on the same SAN is possible.

When configuring the storage on the MSA1000, two methods are available: the CLI is available online, and the ACU can be run from the MSA1000 Support Software CD.

Supported Operating Systems

- NetWare 5.1
- NetWare 6.0

Note: For a current listing of the supported operating system versions, see the MSA1000 website www.hp.com/go/msa1000.

Supported Software Components

- Array Configuration Utility
- Command Line Utility
- Insight Manager 7
- HP StorageWorks Secure Path for NetWare
- Storage Event Broker for NetWare
- SmartStart and Support Software version 5.50 or later

Supported Host Bus Adapters

- FCA2210 Host Bus Adapter

Specific Configuration Needs

Initial system configuration tasks are detailed in the NetWare section of the *Modular SAN Array 1000 Quick Installation Guide*. This guide is included in the Modular SAN Array 1000 Setup and Management kit.

OpenVMS Environments

OpenVMS configurations that use the MSA1000 may range from a simple single-controller MSA1000 connected to a single node to a complex fabric with redundant-controller MSA1000s supporting multiple clusters with many nodes.

Supported Operating System Versions

For Fabric topology Fibre Channel SANs, the following OpenVMS operating systems are currently supported:

- OpenVMS Alpha Version 7.2-2 or later

For Loop topology Fibre Channel SANs, the following OpenVMS operating system is currently supported:

- OpenVMS Alpha Version 7.3-1 or later

Note: For a current listing of the supported operating systems, see the MSA1000 website at www.hp.com/go/msa1000.

Note: If the MSA1000 is deployed with redundant controllers, you must download and install a FIBRE_SCSI remedial kit to access the storage of the MSA1000.You can download the remedial kits from the HP website or contact your support representative to receive the remedial kits on media appropriate for your system.

Supported Software Components

- Command Line Interface (CLI)
- Compaq Analyze
- Controller Firmware Upgrade Utility

Supported Host Bus Adapters

- FCA2354 Host Bus Adapter

Specific Configuration Needs

- The Modular SAN Array 1000 Setup and Management Kit is shipped with the MSA1000, but is not applicable for OpenVMS deployments. This kit is used by Windows, Linux, and NetWare operating system environments.
- To set up your OpenVMS system to recognize the MSA1000, you must obtain the proper version of the operating system. Use your normal support channels to obtain and load the new operating system.
- If you plan to operate the MSA1000 in a redundant-controller configuration, you must first download and install a FIBRE_SCSI remedial kit for your operating system.

  FIBRE_SCSI remedial kits are available for OpenVMS Alpha 7.3-1, 7.3, and 7.2-2 operating system versions.

  You can download the remedial kits from the HP website or contact your support representative to receive the remedial kits on media appropriate for your system.

- OpenVMS supports access to MSA1000 storage connected to Fabric topology Fibre Channel SANs and Loop topology Fibre Channel SANs. See the HP website for a current listing of the supported operating system versions for these different configuration types.
- Each controller and storage unit must be assigned a unique ID number. These unique numbers identify each device to the operating system. The range of valid ID numbers is 1 through 9999.

  Each ID number must be unique to the type of device in the entire SAN that is visible to the OpenVMS system. Data loss or corruption can occur from duplicate ID numbers. Although a controller and a storage unit may share the same ID, duplicate controllers must have unique IDs and all storage units must have unique IDs (even when they are managed by different controllers).

  The controller ID is used to create a device name in the form *1GGAnnnn*, where *nnnn* is the ID number.

  The storage unit ID is used to create a device name in the form *1DGAnnnn*, where *nnnn* is the ID number.

  For information on setting the ID number associated with a controller, see "Array Controller Configuration Commands" in the "Command Line Interface (CLI)" chapter.

  For information on setting the ID number associated with each storage unit, see "LUN Management Commands" in the "Command Line Interface (CLI)" chapter.
- To access storage units or controllers that are added to the SAN without re-booting the OpenVMS system, issue the following command from the server:

  ```
  $ MCR SYSMAN IO AUTOCONFIGURE
  ```
- For more information on using Fibre Channel devices on OpenVMS, see the *OpenVMS Cluster Configurations* and the *OpenVMS Version 7.3-1 Release Notes*.

Tru64 UNIX Environments

Tru64 UNIX currently supports up to four nodes in a cluster when using Fabric configuration and a two nodes in a cluster when using Arbitrated Loop.

Supported Operating System Versions

- 5.1A with NHD6 and PK4 offers boot and data device support

  (NDHD represents New Hardware Delivery kit)
  (PK represents Patch Kit)
- 5.1B with NHD6 and IPK offers boot and data device support

  (IPK represents Initial Patch Kit)

Note: For a current listing of the supported operating system versions, see the MSA1000 website at www.hp.com/go/msa1000.

Supported Software Components

- Command Line Interface (CLI)
- Compaq Analyze
- Controller Firmware Upgrade Utility

Supported Host Bus Adapters

- FCA2354 Host Bus Adapter

Specific Configuration Needs

- The Modular SAN Array 1000 Setup and Management Kit is shipped with the MSA1000, but is not applicable for Tru64 UNIX deployments. This kit is used by Windows, Linux, and NetWare operating system environments.
- To set up your Tru64 UNIX system to recognize the MSA1000, you must obtain the proper version of the operating system. Use your normal support channels.

MSA1000 Setup and Sample Configurations

3

This chapter includes the following sections:

- MSA1000 Setup, page 34
- MSA1000 Sample Configurations, page 42

MSA1000 Setup

This section is a brief checklist of the setup procedures for the MSA1000. Detailed instructions are not included in this section, but are documented in other chapters of this guide or in additional documents. All references are included. In addition, an overview of the MSA1000 installation procedures can be found in the *HP StorageWorks Modular SAN Array 1000 Installation Overview* document.

Installation and setup of the MSA1000 includes:

- Installing and Connecting Hardware Components and Cables
- Installing Support Software
- Configuring the Switches, Routers, and Servers
- Setting up and Configuring the Storage
- Identifying the Server Host Mode

Each of these steps is discussed in the following sections.

Installing and Connecting Hardware Components and Cables

The physical installation of the MSA1000 includes:

- Installing MSA1000 Option Kits
- Installing the MSA1000 in the Rack
- Installing the Host Bus Adapters in the Servers
- Connecting the Cables

Installing MSA1000 Option Kits

If any modifications to the MSA1000 standard shipping configuration need to be made, it may be easier to make these changes before installing the MSA1000 in its rack. Changes may include installing redundant controllers, additional controller cache, embedded switches, or embedded hubs.

Detailed instructions for replacing and installing MSA1000 components are found in other chapters of this guide or separate documents.

For instructions on installing a redundant controller, refer to the *HP StorageWorks Redundant Controller Option Kit Installation Instructions* packaged with the MSA1000 Controller.

For instructions on upgrading the MSA1000 Controller cache, see "Replacing the MSA1000 Controller Cache" in the "Operation and Management" chapter of this guide.

For instructions on installing an embedded MSA SAN Switch 2/8, refer to the *HP StorageWorks MSA SAN Switch 2/8 Installation Guide* packaged with the switch.

For instructions on installing an embedded MSA Hub 2/3, refer to the *HP StorageWorks MSA 2/3 Hub Installation Instructions* packaged with the hub.

Installing the MSA1000 in the Rack

Depending on the rack, the procedures for installing the MSA1000 will differ. The MSA1000 is shipped with the necessary mounting rails. Refer to the documentation that came with the rack for specific instructions.

Installing the Host Bus Adapters in the Servers

The supported HBAs for the different operating systems are listed in Table 6.

Table 6: Supported HBAs

Operating System	HBA
Windows	FCA2101 Host Bus Adapter 64-bit/33-MHz PCI-to-Fibre Channel Host Bus Adapter
NetWare	FCA2210 Host Bus Adapter
Linux	FCA2214 Host Bus Adapter
OpenVMS	FCA2354 Host Bus Adapter
Tru64 UNIX	FCA2354 Host Bus Adapter

Refer to the instructions that came with the server or the HBA for instructions on installing the HBAs.

Connecting the Cables

Fiber optic cables are used to connect the MSA1000 to interconnect devices, such as hubs and switches, and if additional storage enclosures are being used, SCSI cables are used to connect the MSA1000 to the additional storage enclosures.

Fiber optic cables used will differ, depending on the components and devices being connected. The MSA1000 and all of its embedded components use a 2 Gb/s Lucent connector (LC). Depending on the connector in the HBAs, external switches, or hubs, a 1 Gb/s subscriber connector (SC) or a 2 Gb/s LC connector will be needed on the non-MSA1000 end of the cable.

SCSI cables are used to connect up to two additional external storage enclosures to the MSA1000. Use standard VHDCI SCSI cables of appropriate length.

Because there are two interconnect devices and HBAs in redundant configurations, cable connections must follow exacting installation requirements. See the illustrations later in this chapter for examples.

Fiber Optic Cables

Multi-mode fiber optic cables are capable of supporting distances of 2 m to 500 m (6.56 ft to 1640.42 ft) at 1 Gb and up to 300 m (984.25 ft) at 2 Gb. These cables are for use with Short-wave transceivers only. Several multi-mode fiber optic cable option kits are available from HP. Each kit contains a multi-mode fiber optic cable with either a 1 Gb/s or a 2 Gb/s connector attached to each end.

For cables greater than 15 meters, contact an independent fiber optic cable supplier.

If you use an existing 62.5-micron optical cable, you must obtain a 62.5-micron jumper from an independent source. A 50-micron cable cannot be spliced with a 62.5-micron cable. In addition, these cables support half the maximum allowable distance of the multi-mode cables.

Figure 10 provides an illustration of the different connectors.

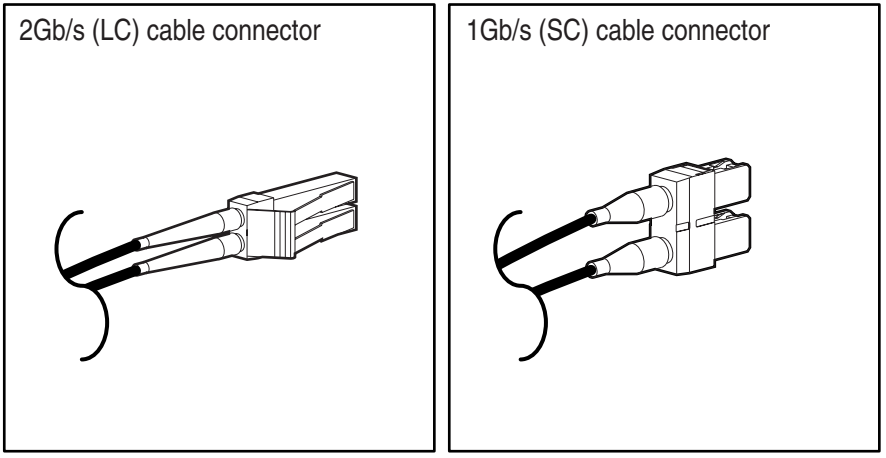

Figure 10: 2 Gb/s and 1 Gb/s connectors

Cable Installation Considerations

To ensure the cabling in the back of a rack system does not interfere with system operation or maintenance, gather the cables in the rear of the MSA1000 and bind them loosely with cable ties. Route the excess cables out of the way, along the side of the rack. See Figure 11 for an illustration.

When the cables are tied together and the excess cable is routed away from the MSA1000, system components and indicators are easily visible and accessible.

Caution: Do not overtighten the cable ties. Damage to the optical cables may result.

Caution: Make certain the fiber optic cables are installed and supported so that no excess weight is placed on the Fibre Channel connectors. This is necessary to prevent damage to the connector and cable and to prevent a cable bend radius less than 3 inches (7.62 cm) at the connector and along the cable length. Excess fiber optic cable should be coiled and tied out of the way, being careful not to coil the cable in a tight loop with a bend radius of less than 3 inches (7.62 cm).

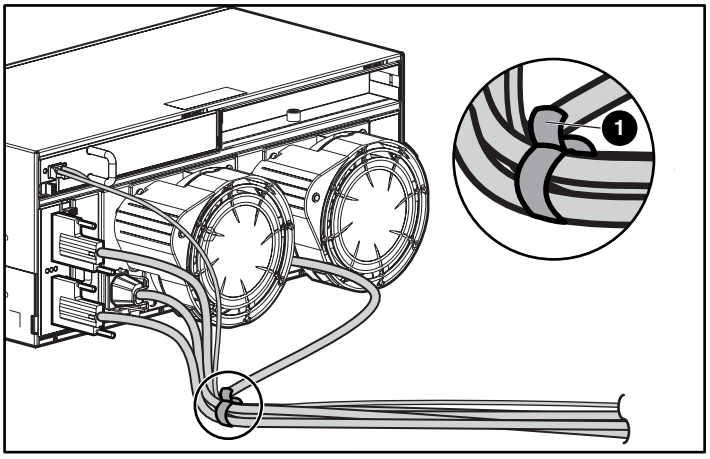

Figure 11: Cable management using cable ties ❶

Installing Support Software

Support software includes HBA drivers as well as the ACU, Insight Manager 7, and Compaq Analyze components that were discussed previously in this chapter.

After the MSA1000 is installed and connected to the server, you can load the needed software components onto the server.

HBA Drivers

All servers that will access the MSA1000 must have the appropriate HBA drivers installed. Drivers for Windows, Linux, and NetWare environments are located on the Modular SAN Array 1000 Support Software CD, shipped with the MSA1000. For instructions on installing these drivers, refer to the *HP StorageWorks Modular SAN Array 1000 Quick Installation Guide* included with the CD in the HP StorageWorks Modular SAN Array 1000 Setup and Management kit.

The HBA drivers for OpenVMS and Tru64 UNIX environments are built into their operating systems. To obtain the latest version of the operating system, use your normal support channels.

Array Configuration Utility

If your operating system supports the ACU and you plan to use the ACU to configure the storage of the MSA1000, it may be loaded from the MSA1000 Software Support CD onto the server. For instructions on installing the ACU, refer to the *HP StorageWorks Modular SAN Array 1000 Quick Installation Guide* included with the CD in the HP StorageWorks Modular SAN Array 1000 Setup and Management kit.

Insight Manager 7

If you are running in a Windows, Linux, or NetWare Linux environment, you can use Insight Manager 7 to monitor the performance of your storage system. Install Insight Manager 7 as instructed in the literature provided with your Insight Manager 7 purchase.

Compaq Analyze

If you are running in an OpenVMS or Tru64 UNIX environment, you can use Compaq Analyze to monitor your storage system. If Compaq Analyze is not already loaded on your server, install it as instructed in the literature provided with Compaq Analyze.

Configuring the Switches, Routers, and Servers

Before the MSA1000 can be accessed, the switches, routers, and servers must be set up. Configuration tasks may include entering the IP addresses and other parameters, setting up zoning on the switch, establishing a workgroup or domain environment, and establishing cascading patterns.

For information on configuring the embedded MSA SAN Switch 2/8, refer to the *HP StorageWorks MSA SAN Switch 2/8 Installation Guide*.

For information on setting up external switches and routers as well as the servers, refer to their respective user manuals.

Setting up and Configuring the Storage

After the servers and interconnect devices are set up and the MSA1000 is physically installed, the storage of the MSA1000 can be accessed and configured.

Storage configuration steps include:

- Determining a Storage Plan
- Configuring the Storage Arrays and Logical Volumes
- Setting up and Configuring Secure Path in Redundant Configurations

Determining a Storage Plan

Planning the storage configuration is critical to the successful and efficient deployment of the MSA1000. Plans should include decisions about total system capacity, fault-tolerance and striping methods of the arrays, sizes of the arrays, and plans for future growth.

Configuring the Storage Arrays and Logical Volumes

After a storage plan is developed, configure the storage by using the available user interfaces of the MSA1000. Depending on the server operating system and user preference, the available user interfaces include the ACU and the CLI.

Table 7: Configuration Tools

Operating System	Configuration User Interface
Windows	ACU, CLI
Linux	ACU, CLI
NetWare	ACU (offline), CLI
OpenVMS	CLI
Tru64 UNIX	CLI

The ACU and the CLI user interfaces are documented in separate chapters of this guide.

For information on using the ACU, see the "Array Configuration Utility (ACU)" chapter.

For information on using the CLI, see the "Command Line Interface (CLI)" chapter.

Setting up and Configuring Secure Path in Redundant Configurations

For Windows, Linux, and NetWare operating systems, redundant configurations with duplicate array controllers, interconnect devices, and HBAs must use HP StorageWorks Secure Path software to ensure continuous data availability.

Secure Path eliminates the hard drive, array controller, interconnect hardware, and HBA as single points of failure in the storage system. These redundant physical connections define separate physical paths. If any component in the path between the host and the storage fails, Secure Path ensures availability by redirecting pending and subsequent requests to an alternate path.

For information on using Secure Path, refer to the documentation provided with your Secure Path software.

Note: Secure Path is not needed or supported for OpenVMS and Tru64 UNIX operating systems, because multipath capability is built into these operating systems.

Identifying the Server Host Mode

For all operating system environments, after the arrays and logical volumes have been created, each server's operating system must be identified to the MSA1000.

The Host Mode can be set in both the ACU and the CLI.

For information on setting the Host Mode using the ACU, see the "Array Configuration Utility (ACU)" chapter.

For information on setting the Host Mode using the CLI, see "Managing Connection Profiles" in the "Command Line Interface (CLI)" chapter.

MSA1000 Sample Configurations

Because the MSA1000 supports a variety of embedded and external switches and hubs in standalone and clustered deployments, many configuration options exist. This section illustrates possible configurations using the following interconnect options:

- Fibre Channel I/O Module Configurations
- Embedded MSA SAN Switch 2/8 Configurations
- Embedded MSA Hub 2/3 Configurations

Note: Redundancy and cluster support is possible in all configurations.

Redundant configurations are used to ensure uninterrupted availability of data. Complete duplication of all hardware components is a requirement of a redundant configuration. This duplicity includes using two MSA controllers, and at least two interconnect devices (depending on the configuration this may include MSA Fibre Channel I/O Modules, MSA SAN Switch 2/8s, MSA Hub 2/3s, or external switches or hubs), and HBAs in the server. As mentioned previously, the fiber cables must be connected in a precise manner to provide the multiple paths required when rerouting data.

The following illustrations provide samples of possible configurations, and are in no way comprehensive.

Fibre Channel I/O Module Configurations

When using the included Fibre Channel I/O Module, the MSA1000 will usually be connected to an external switch or hub. In Windows, Linux, and NetWare environments only, the MSA1000 can be connected directly to a server.

Fiber optic cables are used to connect the MSA1000 to the server and SCSI cables are used to connect the MSA1000 to the additional storage enclosures.

Figure 12 illustrates this basic configuration of connecting one server directly to the Fibre Channel I/O Module of the MSA1000.

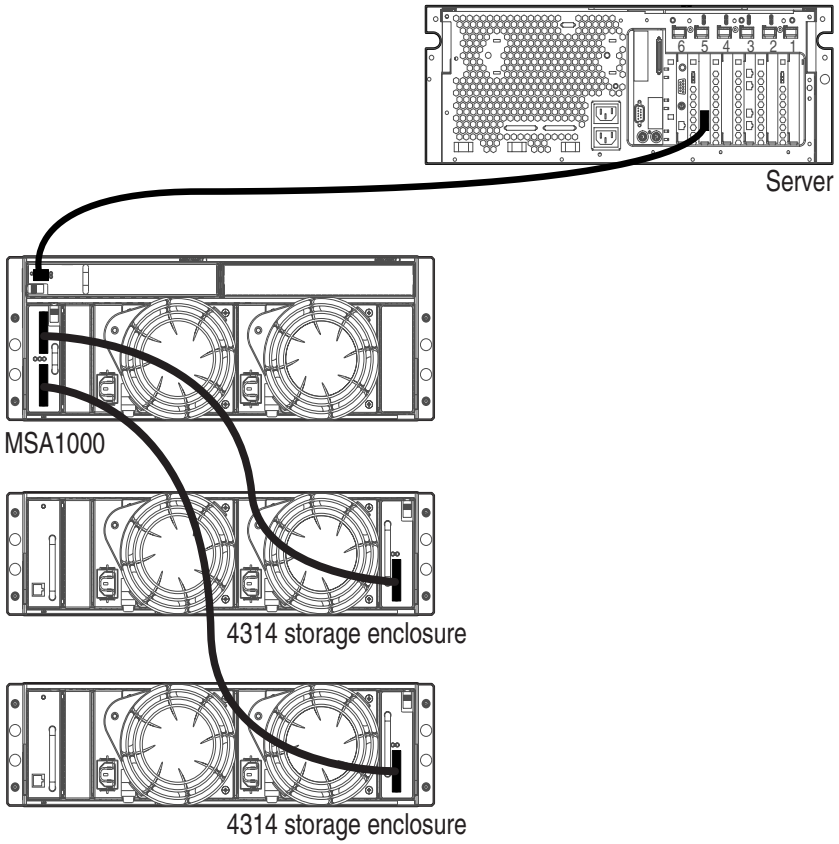

Figure 12: I/O module directly connected to a server

Another common configuration includes connecting an external SAN switch to the Fibre Channel I/O module of the MSA1000.

Figure 13 illustrates multiple servers and multiple MSA1000s connected to a SAN switch in a non-redundant configuration.

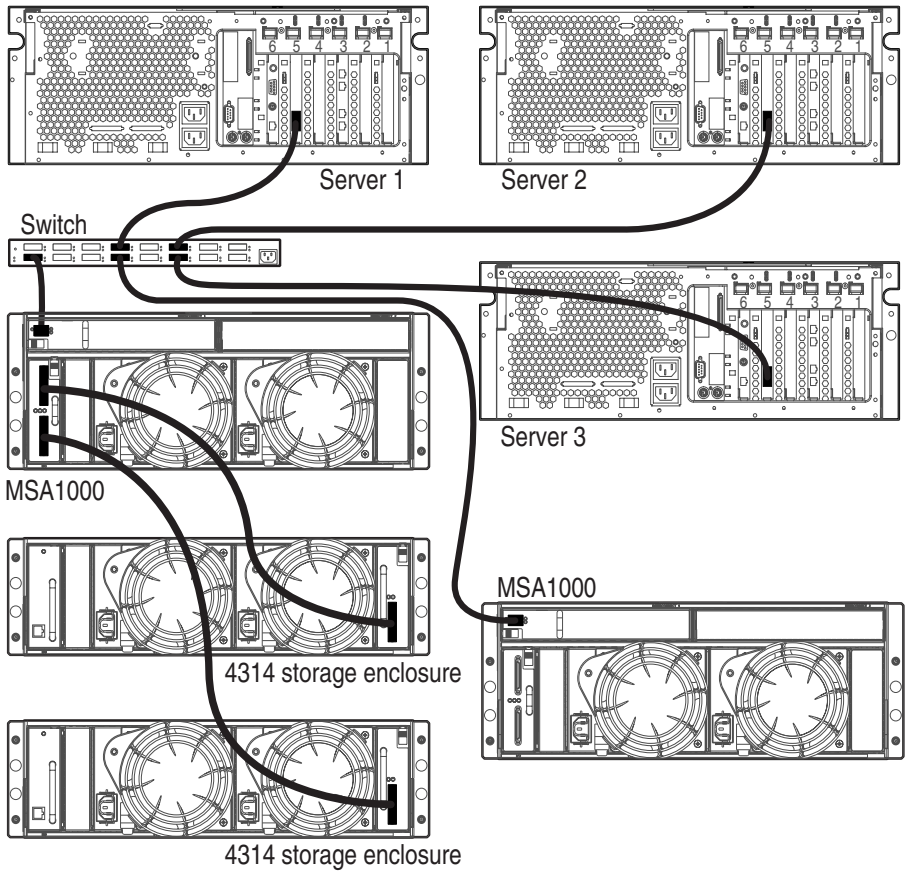

Figure 13: I/O module connected to an external switch, non-redundant configuration

When a redundant configuration is set up using external SAN switches, two MSA1000 Controllers, two Fibre Channel I/O modules, two SAN switches, and two HBAs in each server are required.

Figure 14 illustrates the cable connections of a redundant configuration when external SAN switches are connected to the Fibre Channel I/O modules of the MSA1000.

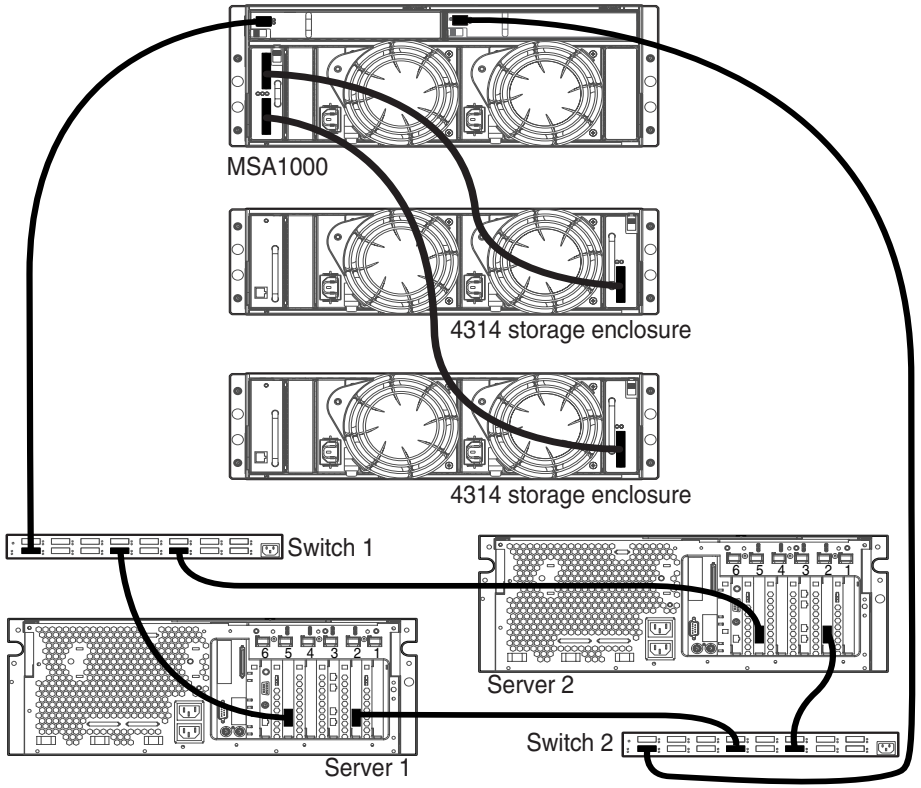

Figure 14: I/O module connected to an external SAN switch, redundant configuration

Embedded MSA SAN Switch 2/8 Configurations

The following illustrations provide examples of an MSA SAN Switch 2/8 installed in the MSA1000.

Figure 15 shows the MSA1000 in a non-redundant configuration and Figure 16 shows the MSA1000 in a redundant configuration.

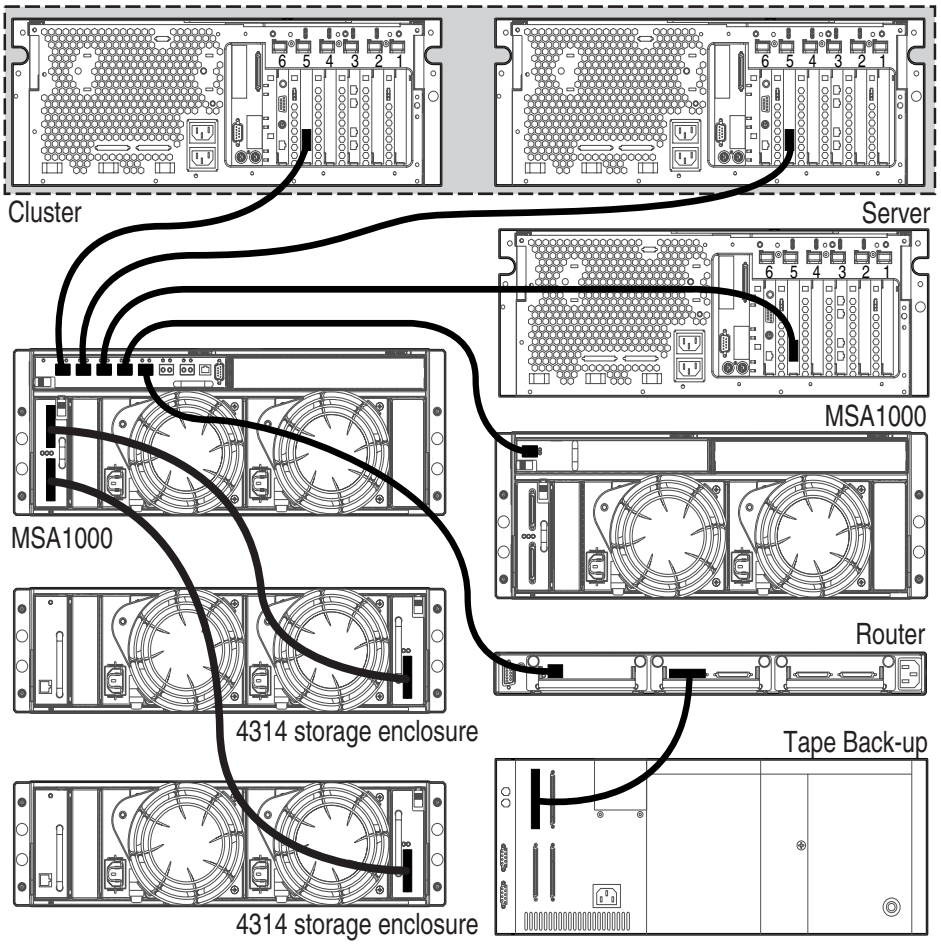

Figure 15: MSA SAN Switch 2/8, non-redundant configuration

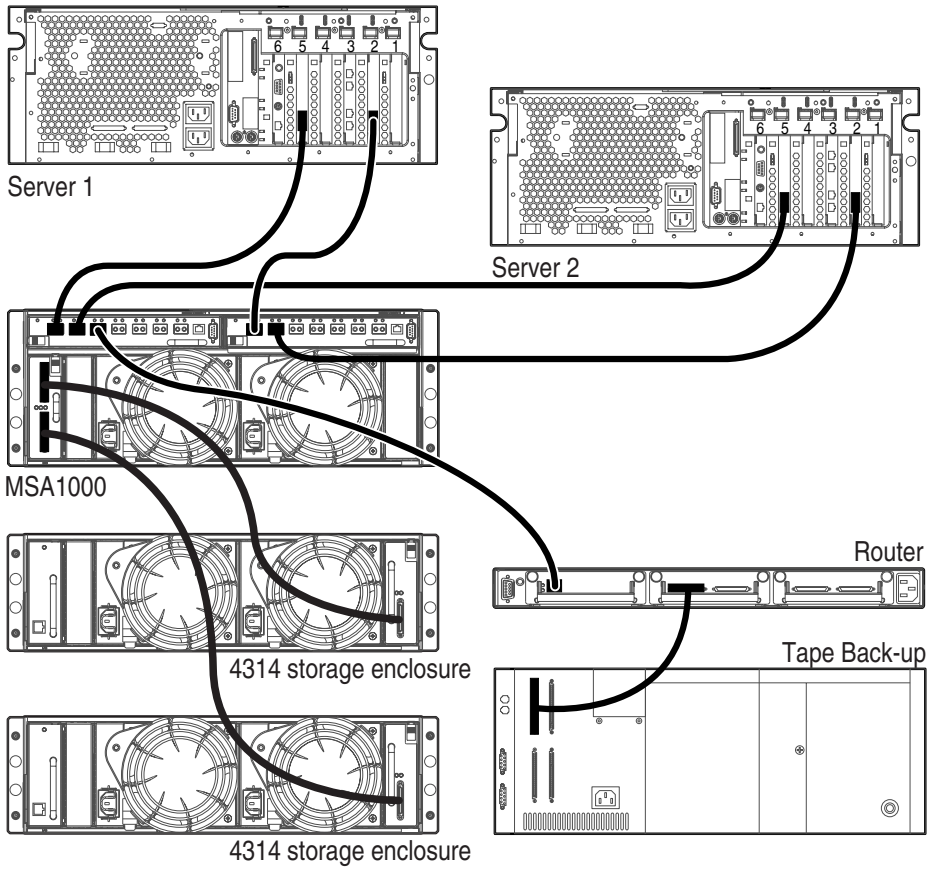

Figure 16: MSA SAN Switch 2/8, redundant configuration

Embedded MSA Hub 2/3 Configurations

Another basic configuration of the MSA1000 includes using the MSA Hub 2/3 option to connect two servers to the storage in a redundant, possibly clustered environment.

Figure 17 is an illustration of this redundant configuration using the embedded hub.

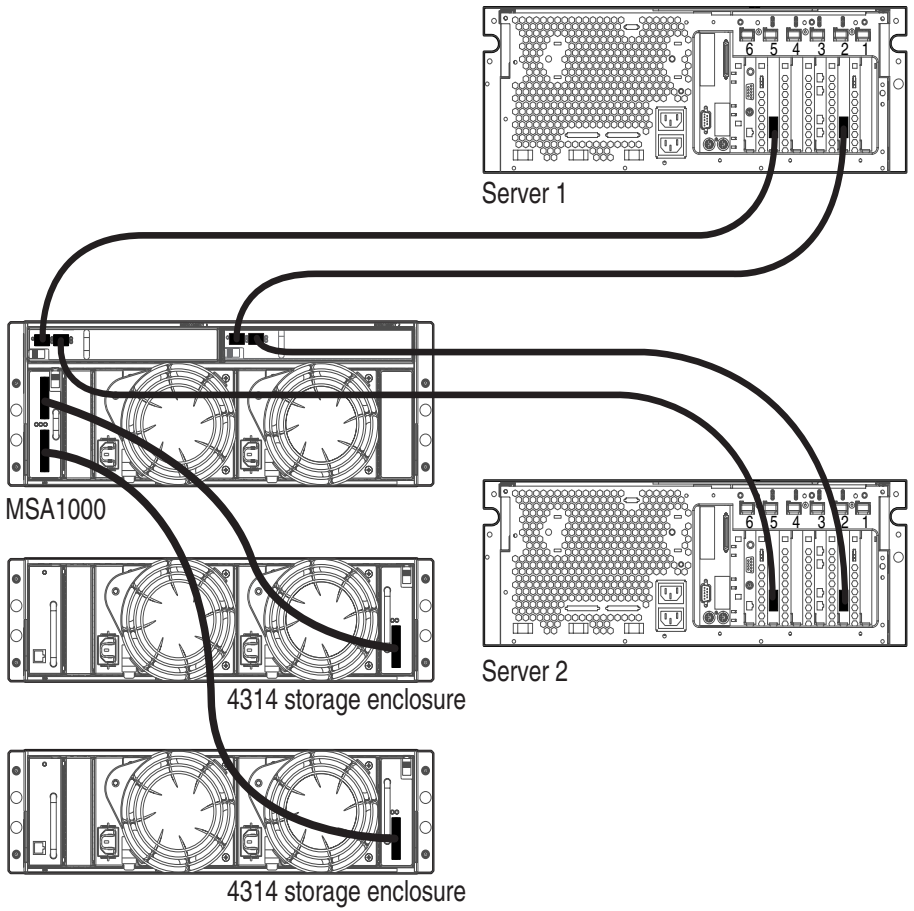

Figure 17: MSA Hub 2/3 connected to two servers, redundant configuration

Operation and Management

4

This chapter describes the operational tasks of managing the MSA1000.

Included sections:

- Connecting the Power, page 50
- Applying Power, page 51
- Interpreting Component Indicators, page 52
- Replacing Components, page 61
- Adding Additional Storage Enclosures, page 79
- Updating the Controller Firmware, page 82

Connecting the Power

The power cord should be approved for use in your country. It must be rated for the product and for the voltage and current marked on the product's electrical ratings label. The voltage and current rating of the cord should be greater than the voltage and current rating marked on the product. In addition, the diameter of the wire must be a minimum of 1.02 mm² or 18 AWG. If you are using 18 AWG, your maximum length may be up to 12 feet.

A power cord should be routed so that it is not likely to be walked on or pinched by items placed upon it or against it. Particular attention should be paid to the plug, electrical outlet, and the point where the cord exits from the product.

After all hardware components are installed and the unit is in place, the power can be connected.

WARNING: To reduce the risk of electric shock or damage to the equipment:

- Do not disable the power cord's grounding plug. The grounding plug is an important safety feature.
- Plug the power cord into a grounded (earthed) electrical outlet that is easily accessible at all times.
- Disconnect power from the storage system by unplugging all power cords from either the electrical outlet or storage system.

1. Plug the AC power cord into a grounded outlet.
2. Plug the AC power cord into the MSA1000.

   The power supply automatically senses the input voltage. It is not necessary to select the correct main voltage.
3. Plug the second AC power cord into the redundant power supply.
4. Plug the second AC power cord into a grounded outlet nearby.

Applying Power

Before applying power to the MSA1000, all components of the storage system must be installed and connected to the supported interconnect options. Hard drives should be installed in the MSA1000 so that they can be identified and configured at power up.

The MSA1000 components should be powered up in the following sequence:

1. **Storage Expansion Enclosures**
2. **MSA1000**

   Power on the MSA1000 using the switch located in the far-right lower area of the front panel of the MSA1000. After powering on, wait until the message "`MSA1000 Startup Complete`" appears on your display. This process may take up to two minutes.
3. **Interconnect Devices**
4. **Server(s)**

Interpreting Component Indicators

The MSA1000 and each of its components are equipped with a series of indicators.

If the fault indicators on any of the MSA1000 components are amber or red, or if Insight Manager or Compaq Analyze indicates a fault, determine the reason for this alert immediately. Examine the component indicators to see if any indicates a fault.

The following sections describe the indicators:

- MSA1000 Controller Indicators
- Enclosure Status Indicators
- Power Supply/Blower Assembly Indicators
- SCSI I/O Module with Integrated EMU Indicators
- Hard Drive Indicators

MSA1000 Controller Indicators

During normal runtime, the MSA1000 Controller has 18 indicators that indicate activity or malfunction of the controller. They are labeled 0-17, as shown in the following figure. The table that follows describes the purpose and function of each indicator.

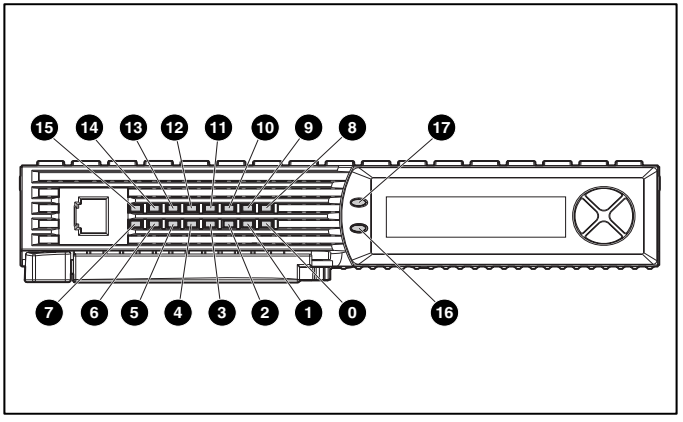

Figure 18: MSA1000 Controller indicators

Table 8: MSA1000 Controller Indicator Descriptions

Indicator	Function	Description
❶-❷	Busy status	These three LEDs are used to progressively represent the processing load on the controller. ON = the controller is idle. OFF = the controller is operating at full capacity.
❸-❼	Fibre Channel ID	Indicates the 5-bit Arbitrated Loop Physical Address (ALPA) assigned to this array controller (not applicable when in fabric mode).
❽	Idle Heartbeat	Indicates the controller is idle and functioning.
❾	Active/Standby	ON=Controller is active. OFF=Controller is in standby.
❿	Direct Memory Access (DMA) active	ON = DMA transfers are active.
⓫	Logical I/O active	ON = Currently processing logical requests from the Host Adapter.
⓬	SCSI Bus 0 active	ON = Indicates requests are outstanding on the first SCSI bus.
⓭	SCSI Bus 1 active	ON = Indicates requests are outstanding on the second SCSI bus.
⓮	Cache Activity	ON = Cache active. OFF = No cache activity. Blinking = Cache transfer pending.
⓯	Drive failure	ON = A configured hard drive has failed in the array.
⓰	Redundancy Active	Green indicates two controllers are in a redundant mode of operation.
⓱	Fault	Red indicates an error message has been sent to the controller display.

Enclosure Status Indicators

The front panel indicators are found on the MSA1000, as illustrated in Figure 19.

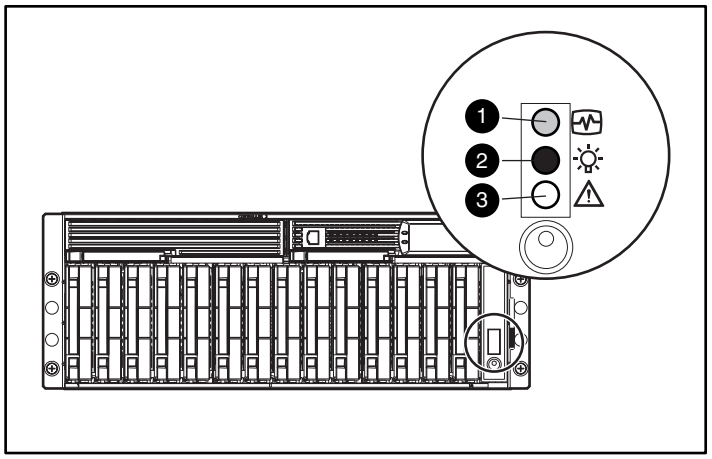

Figure 19: Enclosure status indicators

WARNING: The system power in the MSA1000 does not shut off completely with the power switch. STANDBY removes power from most of the electronics and the drives, but portions of the power supply and some internal circuitry remain active. To remove the power completely, disconnect all power cords from the equipment.

Number	Indicator	Condition	Meaning
❶	EMU	Blinking Green	EMU is operating normally.
		Off or On without blinking	EMU is not operating normally.
❷	Power	Green	System power ON.
		Off	System in STANDBY or system power has been removed.
❸	Fault	Amber	Fault detected in one or more subsystems.
		Off	No faults detected.

Power Supply/Blower Assembly Indicators

The green indicators on the power supply/blower assembly are on when both the power supply and the blower are operational. When a power supply or blower fault occurs, the power supply indicators go off. If the power supply indicators are off, AC power is not present or there is a complete power supply failure.

The power supply-mounted blowers cool the enclosure by circulating air through the enclosure and elements. The rate at which air moves (the airflow) determines the amount of cooling. This airflow is a function of blower speed (rpm). These blowers, under the control of the Environmental Monitoring Unit (EMU) or the associated power supply, can operate at multiple speeds. This ensures that when the enclosure temperature changes, the blowers can automatically adjust the airflow.

If one blower operates too slowly or completely shuts off, the other blower will begin to operate at a higher speed. At the same time, the error condition is reported to LCD panel on the front of the Controller.

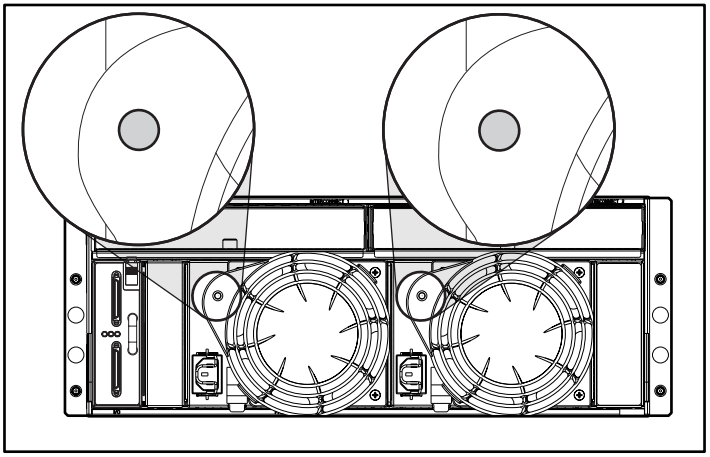

Figure 20: Power supply/blower assembly indicators

SCSI I/O Module with Integrated EMU Indicators

The SCSI I/O Module has three indicators. The following figure and table describe their locations and meanings.

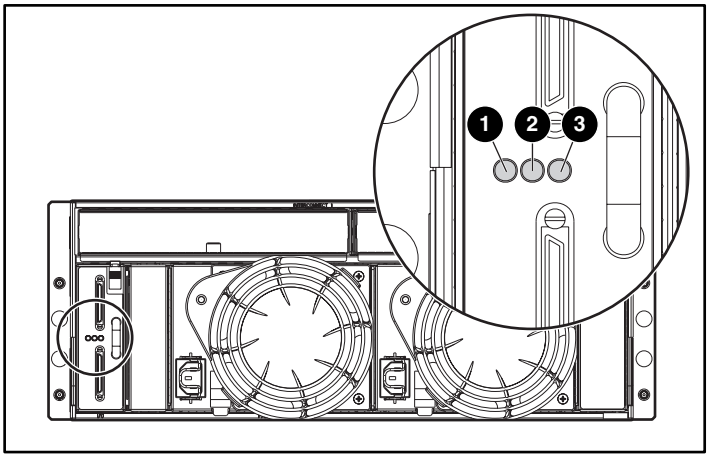

Figure 21: SCSI I/O Module with Integrated EMU

Number	Indicator	Condition	Meaning
❶	Power	Solid green	Power on
		Off	Power off
❷	SCSI Host Port A	Green/Flashing	On/Activity
		Off	Off
❸	SCSI Host Port B	Green/Flashing	On/Activity
		Off	Off

Hard Drive Indicators

The three disk drive status indicators displayed in the following figure define the operational status of each drive.

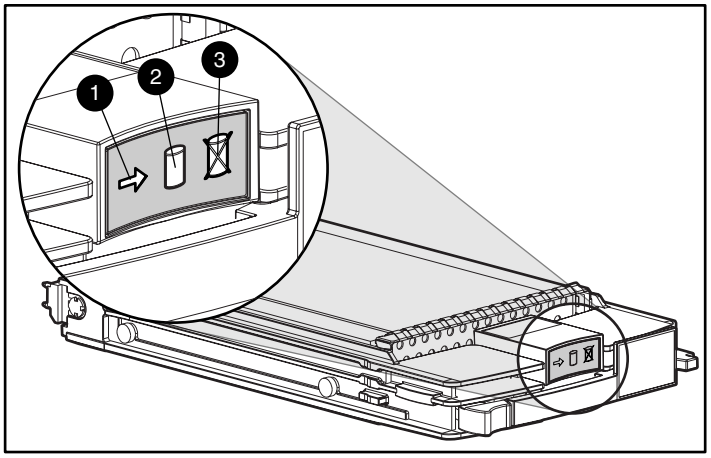

Figure 22: Hard drive indicators

Look for red Drive Failure indicators on any hot-pluggable drive. If any Drive Failure indicator appears amber, replace that drive as soon as possible. Refer to the section "Replacing Hard Drives," later in this chapter, for important information on when it is safe to replace drives.

Table 9: Determining Hard Drive Status from Indicator Combinations

Activity ❶	Online ❷	Fault ❸	Meaning
On	Off	Off	The drive is being accessed, but it is: ■ Not configured as part of an array. ■ A replacement drive and rebuild has not yet started. ■ Spinning up during POST.
Flashing	Flashing	Flashing	**Do not remove the drive. Removing a drive during this process can cause data loss in non-fault-tolerant configurations.** The drive is part of an array being selected by the Array Configuration Utility.

Table 9: Determining Hard Drive Status from Indicator Combinations

Activity ❶	Online ❷	Fault ❸	Meaning
Off	Off	Off	OK to replace the drive online if: The drive is not configured as part of an array, *or* The drive is configured as an online spare.
On	Flashing	Off	**Do not remove the drive. Removing a drive during this process can cause data loss in non-fault-tolerant configurations** The drive is a replacement drive and is being rebuilt *or* If all online lights in a drive array are blinking, an expansion is occurring.
Off	Off	On	OK to replace the drive online. The drive has failed and has been placed offline.
Off, On, or flashing	On	Off	**Do not remove the drive. Removing a drive during this process can cause data loss in non-fault-tolerant configurations.** The drive is online and configured as part of an array.
Off, On, or flashing	On or Off	Flashing	■ A predictive failure alert has been received for this drive. Replace the drive as soon as possible. OR ■ A "locate" command is being executed and the system is identifying the requested drive.

Fibre Channel I/O Module Indicators

The Fibre Channel I/O Module has three indicator lights, as illustrated in Figure 23.

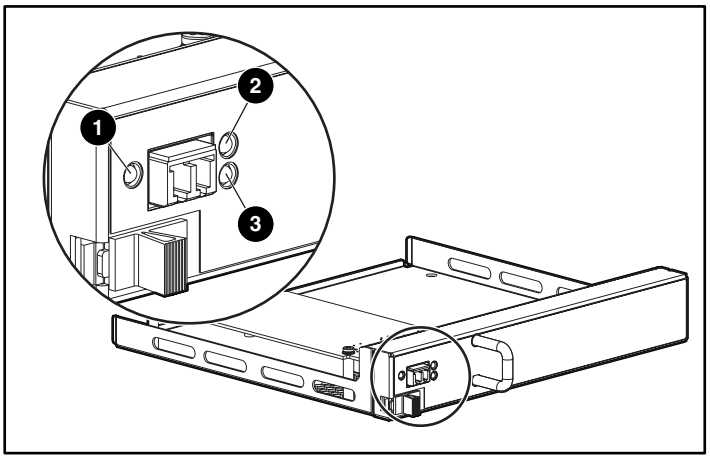

Figure 23: Fibre Channel I/O module indicators

Item ❶ is the Global Status indicator and items ❷ and ❸ are the Link Status Indicators. If the controller is not installed, the two link status indicators will flash amber.

The Global Status indicator and the Link Status indicators are described in separate tables.

Table 10 defines different illuminations of the Global Status Indicator ❶.

Table 10: Fibre Channel I/O Module Global Status Indicator

Condition	Description
Off	Power not applied to Fibre I/O module, no Fibre I/O module installed, controller not plugged in, or unsuccessful controller POST.
Solid Green	Power applied and connection established between Fibre I/O module and Fibre Channel controller.
Flashing Amber	Power is applied but Fibre Channel controller is not present.

Table 11 defines the illuminations of the Link Status Indicators ❷ and ❸.

Table 11: Fibre Channel I/O Module Link Status Indicators

1-Gb Indicator ❶	2-Gb Indicator ❷	Description
Off	Off	A link has not been established since the application of power, initiation of reset, hot-plug of Fibre I/O module, the removal of transceiver or the hot-plug of the transceiver.
Green	Off	A 1-Gb link is active and the status is good.
Flashing Amber	Off	A 1-Gb link was established but is now inactive.
Off	Green	A 2-Gb link is active and the status is good.
Off	Flashing Amber	A 2-Gb link was established but is now inactive.
Flashing Amber	Flashing Amber	The controller was removed or there is no Fibre Card installed.

Replacing Components

Caution: It is important to follow these instructions when replacing components in the MSA1000. If the procedure is done improperly, it is possible to lose data or damage equipment. Refer to Appendix B, "Electrostatic Discharge," for important information on using the proper procedures.

This section includes:

- Replacing the MSA1000 Controller
- Replacing the MSA1000 Controller Cache
- Replacing the Controller Cache Battery Pack
- Replacing a Variable Speed Blower
- Replacing a Power Supply
- Replacing Hard Drives
- Replacing the 2-Gb Small Form Factor Pluggable (SFP) Transceiver

Note: To take advantage of all features, check the MSA1000 website at www.hp.com/go/msa1000, to make sure you have the latest firmware.

Replacing the MSA1000 Controller

If an MSA1000 Controller is failing, informational or error messages will be displayed on the LCD panel of that controller, depending on the condition.

The following steps detail how to replace a failed MSA1000 Controller, but can also be used to re-seat the Controller.

Note: Redundancy is supported during a hard drive expansion, migration, or extension process and during regular drive rebuilds.

Note: Replacement MSA1000 Controllers include a new cache module. Remove this new module from the replacement MSA1000 Controller and replace it with the cache module from the failed controller. Using the same cache module will complete the disk writes that may have been trapped in the Controller's cache. See "Replacing the MSA1000 Controller Cache" for procedural instructions.

1. Press the controller's thumb latch and pull the latch handle toward you ❶.
2. Remove the MSA1000 Controller by pulling it straight out of the chassis ❷, as illustrated in Figure 24.

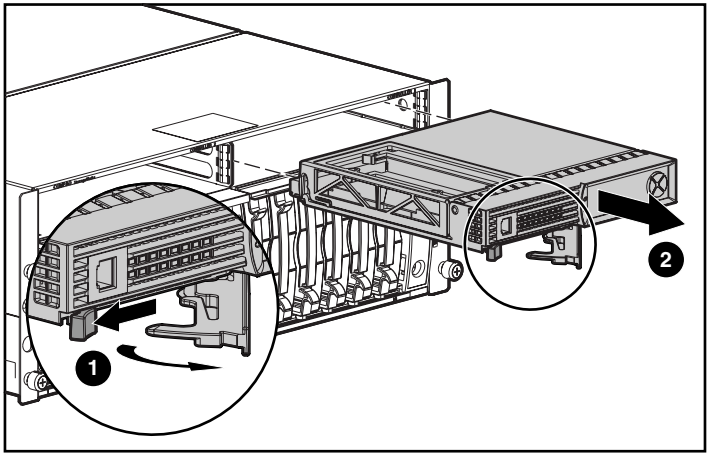

Figure 24: Removing the MSA1000 Controller

3. As illustrated in Figure 25, insert the replacement controller into the chassis ❶.

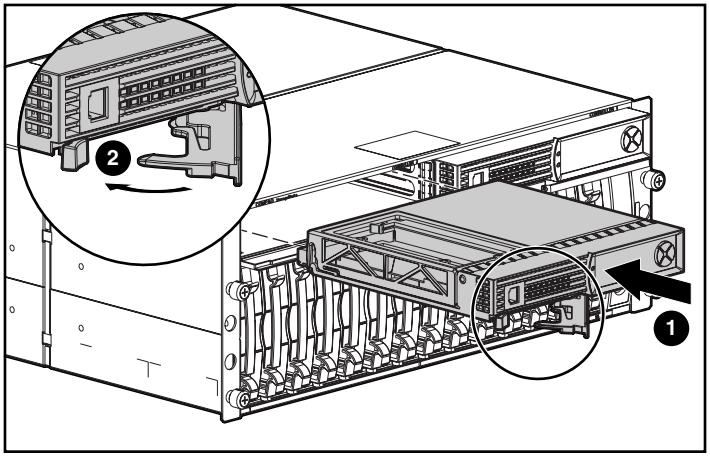

Figure 25: Installing the replacement controller

4. Push the controller in as far as it will go; press the latch inward until it is flush against the front panel ❷.

Replacing the MSA1000 Controller Cache

Caution: It is important to follow these instructions when replacing components in the MSA1000. If the procedure is done improperly, it is possible to lose data or damage equipment. Refer to Appendix B, "Electrostatic Discharge," for important information on using the proper procedures.

Note: If your system is equipped with a single controller, and you must replace the controller cache, you must power down the system first. If your system is equipped with two controllers, and you want to replace a failed cache module with another of the same size, you can replace the module while the system is running. If your system is equipped with two controllers, and you are replacing the cache module with a module of a different size, you must power down the system first, and then change the cache module on both controllers at the same time.

1. Press the controller's thumb latch and pull the latch handle towards you ❶. See Figure 26.
2. Remove the MSA1000 Controller by pulling it straight out of the chassis ❷.

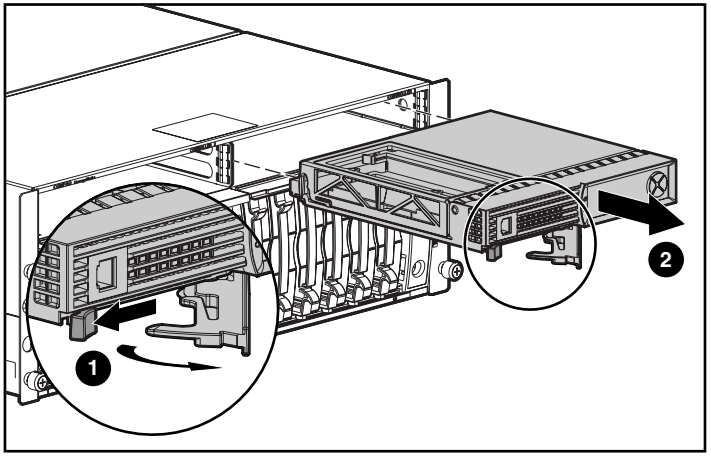

Figure 26: Removing the controller

3. As illustrated in Figure 27, unlatch the clips that are holding the MSA1000 Controller cover ❶ and raise the cover ❷.
4. Simultaneously unlatch the clips that are holding the MSA1000 Controller cache in place ❸.
5. Carefully pull the cache away from the controller board ❹.

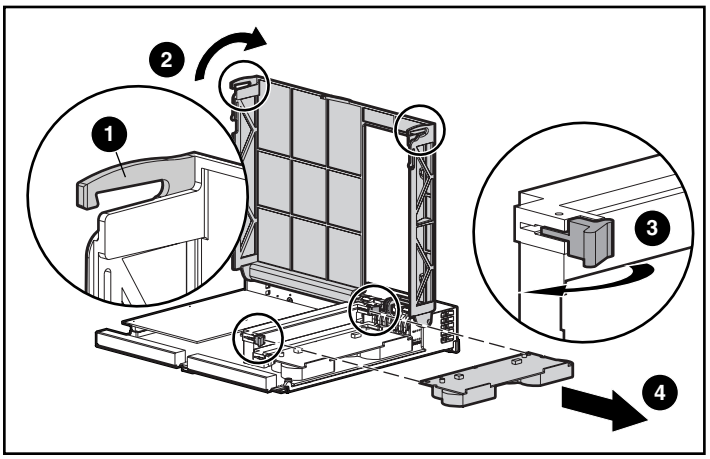

Figure 27: Removing the cache module

6. Install the new MSA1000 Controller cache by sliding the new MSA1000 Controller cache into the controller ❶. Be sure the side latches are fully engaged ❷. See Figure 28.

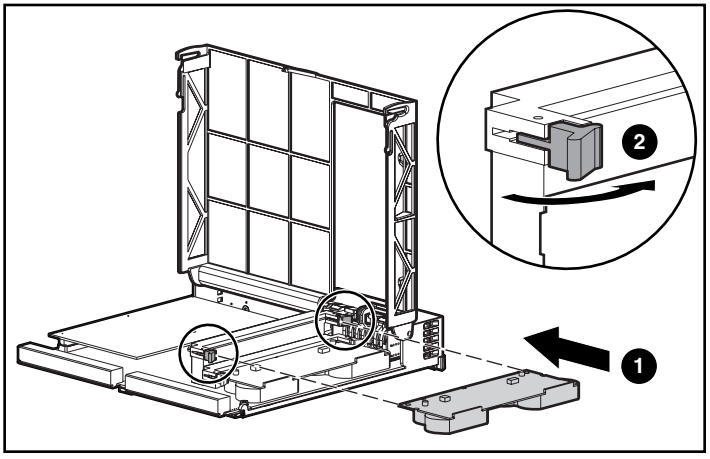

Figure 28: Installing the cache module

7. Push the controller in as far as it will go ❶; press the latch inward until it is flush against the front panel ❷. See Figure 29.

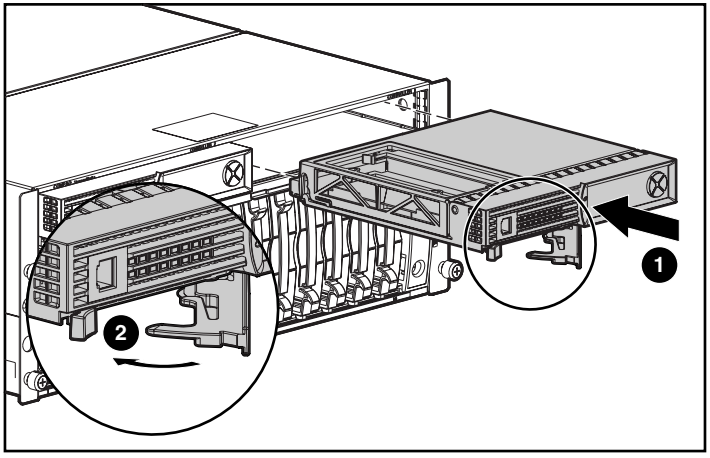

Figure 29: Installing the controller

Replacing the Controller Cache Battery Pack

WARNING: There is a risk of explosion, fire, or personal injury if the battery pack is replaced incorrectly or mistreated. To reduce the risk:

- Do not attempt to recharge the battery outside of the controller.
- Do not expose to water, or to temperatures higher than 60°C.
- Do not abuse, disassemble, crush, puncture, short external contacts, or dispose of in fire or water.
- Replace only with the spare designated for this product.
- Array Accelerator battery disposal should comply with local regulations. Alternatively, return them by established parts return methods to Hewlett- Packard Corporation for disposal.

Caution: It is important to follow these instructions when replacing components in the MSA1000. If the procedure is done improperly, it is possible to lose data or damage equipment. Refer to the "Electrostatic Discharge" appendix for important information on using the proper procedures.

To remove the old NiMH battery pack:

1. Remove the MSA1000 Controller Cache, as instructed in the previous section, "Replacing the MSA1000 Controller Cache."
2. Push down on the bottom clip of the battery pack, attached near the lower corner of the Array Accelerator.

   See Figure 30 for an illustration.

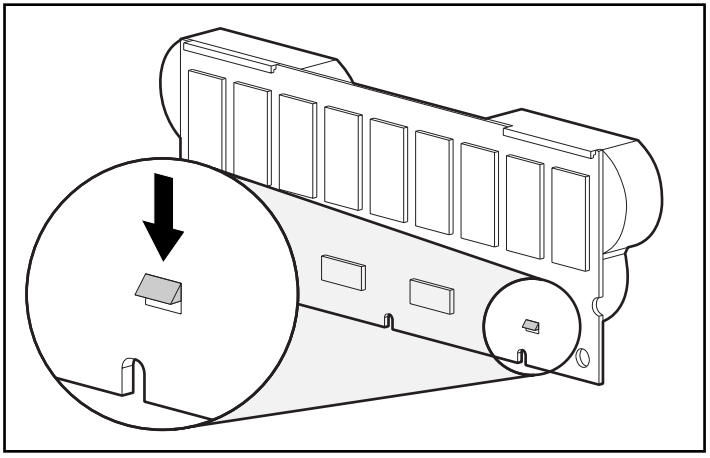

Figure 30: Bottom clip on battery pack

3. Swing the battery pack away from the Array Accelerator to about a 30-degree angle.

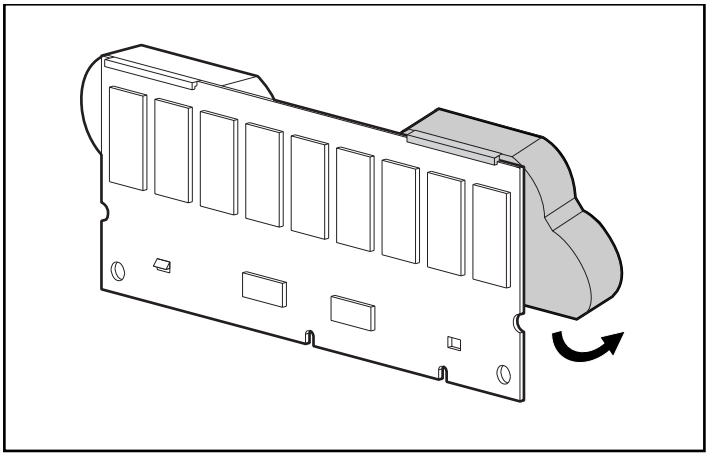

Figure 31: Angling the battery pack

4. Lift the pack upward to unhook the top of the battery pack.

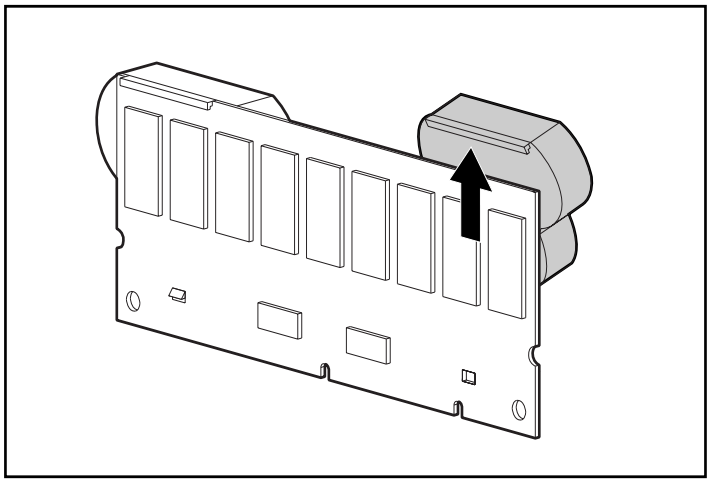

Figure 32: Removing the battery pack

Wait about 15 seconds after removing the old battery packs to allow the battery charge monitor to reset.

Note: Repeat the replacement procedure for any batteries that were installed at the same time as the batteries that were removed.

---

5. Install the new NiMH battery pack by hooking the top of the battery pack to the top of the Array Accelerator with the pack held at a 30-degree angle to the plane of the Array Accelerator board.

   See Figure 33 for an illustration.

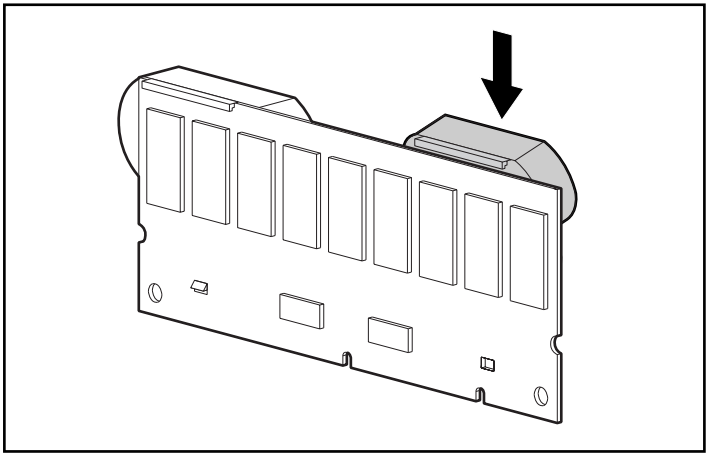

Figure 33: Installing the battery pack

6. After the pack is hooked in position, swing the pack downward making sure the bottom clip and two pegs line up with the holes in the Array Accelerator.
7. Make sure that the top hook ❶ and bottom clip ❷ on the battery pack are securely attached to the Array Accelerator.

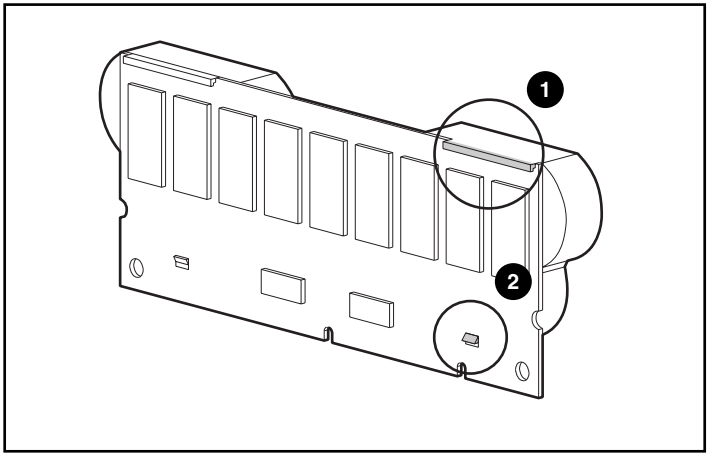

Figure 34: Securing the battery pack

Installation of the new battery pack is complete. Repeat for the other battery pack.

Replacing a Variable Speed Blower

WARNING: The blower blades rotate at a high speed and do not stop *immediately* when power is removed. Allow time for the blades to stop rotating. Avoid touching the rotating blades when removing the blower.

Caution: Removing a power supply significantly changes the airflow within the MSA1000. After removal of a power supply, the remaining supply starts a timer that will shut the system down in five minutes, if the removed power supply is not replaced.

Note: The power supply design ensures that removing a blower does not change the airflow within the enclosure. However, do not remove a blower until the replacement blower is available.

To remove the blower from the power supply:

1. Push in on the two port-colored blower tabs ❶ while pulling on the blower element ❷. See Figure 35 for an illustration.

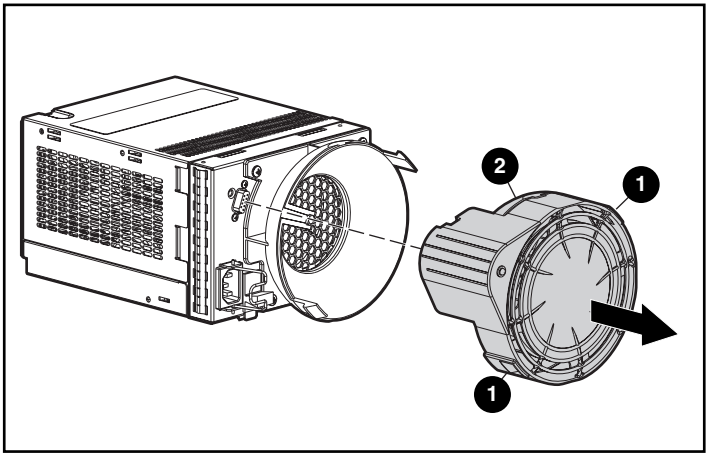

Figure 35: Replacing the blower

Caution: Pressing on the center section of the blower can damage the blades. To prevent this, press only on the outer edge of the blower.

2. Align the new blower's guidepost ❷ with the power supply connector ❶. Slide the replacement blower into the blower base ❹ until the tabs ❸ snap into place.

   See Figure 36 for an illustration.

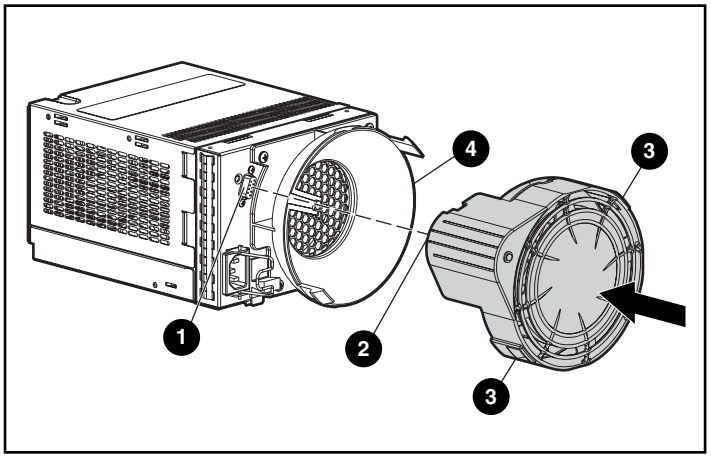

Figure 36: Installing the replacement blower

3. Make sure the following operational blower indications occur:
   a. The blower starts operating immediately.
   b. The blower indicator is On.

Replacing a Power Supply

Caution: Removing a power supply significantly changes the airflow within the enclosure. The system could shut down because of overheating unless the power supply is replaced within *five minute*s.

Note: Replacement power supply assemblies do not include a variable speed blower. You must remove the operational blower from the defective power supply and install the blower on the new supply.

1. Disengage the cord lock ❶ and disconnect the AC power cord from the defective supply.
2. While lifting up on the power supply port-colored module latch ❷, grasp the blower element ❸ and pull the defective power supply out of the enclosure.

Caution: When removing the left power supply, move the cord lock on the right power supply out of the way to avoid dislodging the right power supply and causing your system to become overheated.

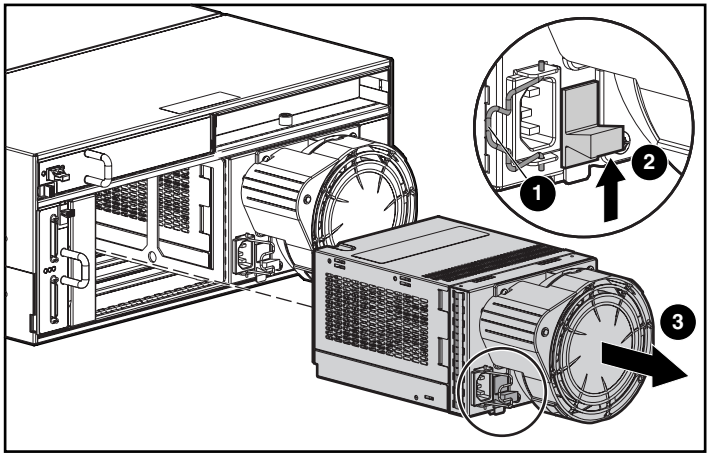

Figure 37: Removing the power supply

3. Install the replacement power supply by lifting up on the power supply module latch ❶ and pushing in on the blower base ❷ until the assembly is fully seated in the enclosure.

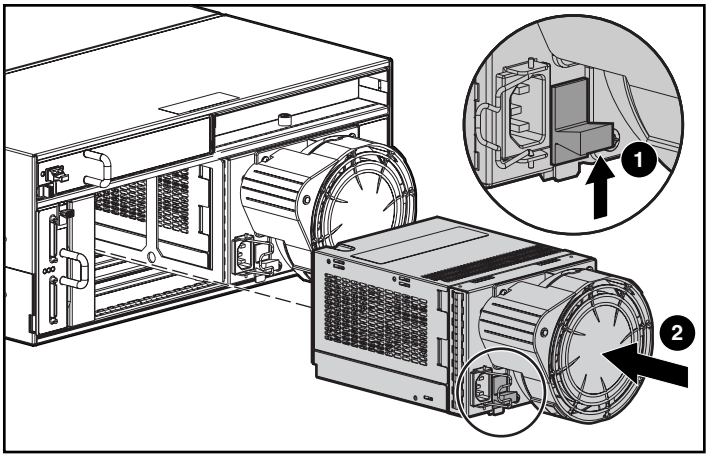

Figure 38: Installing the replacement power supply

4. Re-connect the AC power cord to the replacement power supply and engage the cord lock.

Replacing Hard Drives

Caution: To replace a hot-pluggable drive, it is important to follow the guidelines in this section. Failure to do so could result in data loss and could void your warranty.

Drive failure is indicated by an amber Drive Failure indicator.

RAID 0 is not a fault-tolerant configuration. Never remove a drive from a RAID 0 array unless it has failed. If you are using a RAID 0 configuration, removing an operating drive is not allowed and will result in loss of data.

To remove a working drive without losing data, the entire array must be backed up, the drive must be replaced, and the entire array must be restored. Backing up a single drive and replacing it will not restore the array.

There are some instances in which you may replace a drive in RAID 1, RAID 5, and Advanced Data Guarding configurations. Table 9, located in the earlier section “Hard Drive Indicators,” illustrates the conditions when you may or may not remove a drive in a fault-tolerant system.

Follow these guidelines when replacing drives:

- **Never remove more than one drive at a time (two drives if you are using Advanced Data Guarding)**.

  When you replace a drive, the controller uses data from the other drives in the array to reconstruct data on the replacement drive. If you remove more than one drive, a complete data set is not available to reconstruct data on the replacement drive(s) and permanent data loss could occur.

- **Never remove a working drive**.

  The amber Drive Failure indicator on the drive carrier indicates a drive that has been failed by the controller. Permanent data loss will occur if a working drive is removed while replacing a failed drive.

- **Never remove one drive in an array while another drive in the array is being rebuilt**.

  A drive’s Online indicators flash green (once per second) while it is being rebuilt. A replaced drive is rebuilt from data stored on the other drives.

- **If the system has an online spare drive, wait for it to complete rebuilding before replacing the failed drive**.

  When a drive fails, the online spare becomes active and begins rebuilding as a replacement drive. After the online spare has completed Automatic Data Recovery (the Online indicators will be continuously lit), replace the failed drive with a new replacement drive. Do not replace the failed drive with the online spare. The system will automatically rebuild the replacement drive and reset the spare drive to an available state.

- If you replace a drive while the system is off, it may be necessary to rebuild the replaced drive.

- Up to one half of the drives in a RAID 1+0 array can be replaced at the same time.

To replace a drive:

1. Verify the Activity LED indicator is OFF.

   Hard Drive LEDs are illustrated in Figure 22.

2. Press the ejector ❶ and pivot the lever ❷ to full open position.
3. Pull the drive partially out of the bay and allow time for the disk to stop rotating.
4. Pull the drive completely out of the storage system drive bay ❸.

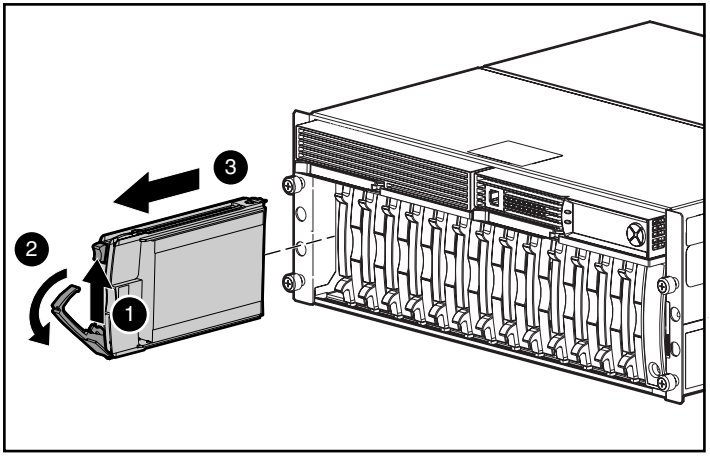

Figure 39: Opening the drive tray latches

5. On the new hard drive, press the ejector and pivot the lever to a full open position.

6. Insert the new drive into the same bay as the drive just removed, sliding it in as far as it will go ❶.

   Make sure the ejector lever ❷ is in the full open position to ensure a correct latch.

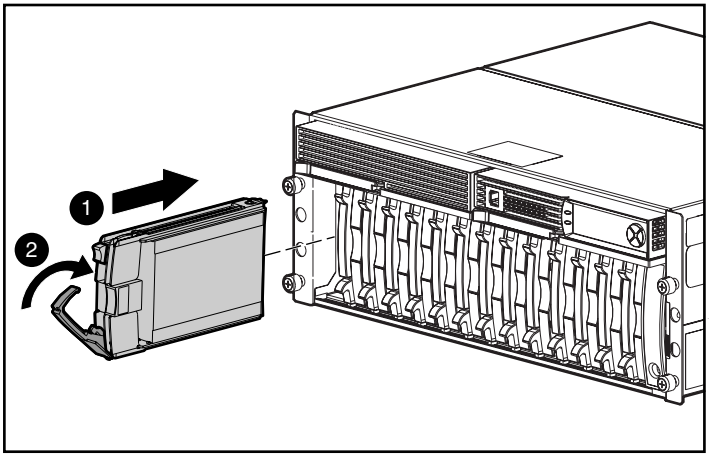

Figure 40: Inserting the new drive

7. Close the ejector lever against the front of the drive. The locking hinge must be engaged behind the front panel of the unit.
8. Insert the drive completely to ensure that the drive is firmly engaged and that the ejector lever is latched.

Caution: Data can be lost if the drive is not firmly seated.

The drive tray indicators light one at a time and then turn off together to indicate that the system has recognized the new drive. In fault-tolerant configurations, allow the replacement drive to be reconstructed automatically with data from the other drives. While reconstruction is in progress, the online indicators flash.

Replacing the 2-Gb Small Form Factor Pluggable (SFP) Transceiver

If a transceiver fails, follow this procedure to replace it. It is not necessary to power down the system.

WARNING: To reduce the risk of injury from laser radiation or damage to the equipment, observe the following precautions:

- Allow only HP Authorized Service Technicians to repair the equipment.
- Do not open any panels, operate controls, make adjustments, or perform procedures to a laser device other than those specified herein.
- Do not stare into laser beam when panels are open.

1. Pull the transceiver out of the device by pulling up and out on the plastic tab.
2. Press the release clip on the bottom of the cable connector ❷ to remove the Fibre Channel cable from the back of the failed transceiver ❸.

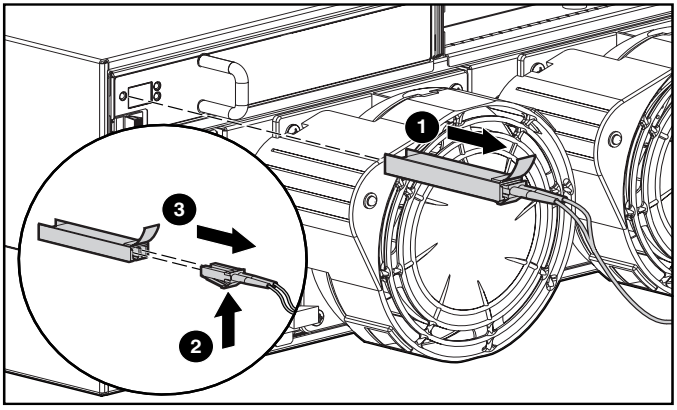

Figure 41: Removing the failed SFP

3. Replace the protective covers on the cable and insert the dust plug covers into the transceiver cage.
4. To insert a new SFP, reverse steps 1 through 3.

Caution: To reduce the risk of damage to the equipment, do not use excessive force when inserting the transceiver.

Adding Additional Storage Enclosures

This section describes how to add HP StorageWorks enclosure models 4314/4214 single bus enclosures or HP StorageWorks enclosure models 4354/4254 dual bus enclosures to the MSA1000 storage subsystem. It does not cover adding disk drives to the MSA1000 or to disk enclosures already attached. Adding disk drives and capacity expansion is discussed in the "Array Configuration Utility" chapter.

Older storage systems with high voltage SCSI connections cannot be attached to the MSA1000. New HP StorageWorks enclosures must be purchased with new Ultra3 Drives.

Methods of expanding:

- Migrating existing storage enclosures from an HP Smart Array to the MSA1000
- Adding a new storage enclosure

Migrating Existing Storage Enclosures

Even though extensive design and testing has been performed, when migrating enclosures it is recommended to perform a back up of your data before you migrate the enclosure to the MSA1000. Perform disk administration steps to defragment your file systems or volumes prior to performing the backup, which allows more efficient use of your backup media and reduces the time for backup.

The following steps provide an overview of migrating a new enclosure to an existing MSA1000:

1. After backing up the storage systems, schedule a convenient time to shut down the servers/applications so the storage subsystem can be turned off.
2. Shut down both the system(s) from which you are removing the disk enclosure(s) and the target system to which you are planning to migrate/consolidate your storage.
3. Disconnect the SCSI cables from your host bus adapter and storage enclosure. If necessary, move the disk enclosure(s) to the new server rack using the existing rack mount hardware.
4. Attach SCSI cables to the MSA1000 expansion SCSI connectors.

   See Figure 42 for an illustration of the SCSI connectors.

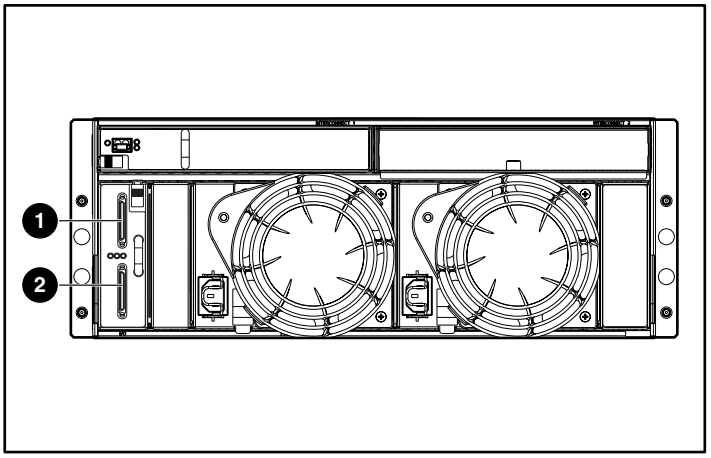

Figure 42: MSA1000 expansion SCSI connectors

Item	Description
❶	SCSI Port A (SCSI Bus 2)
❷	SCSI Port B (SCSI Bus 3)

5. Attach the SCSI cables to the expansion disk enclosures. Attach the power cords to the power supplies at the rear of the enclosures. Power on each external disk enclosure.
6. Power on the MSA1000 system using the power switch in front of the system. Turn on the host server(s) and allow it to boot.
7. Use the ACU or the CLI to verify current configuration is maintained and new volumes are identified. If the migrated volumes are not identified, shut down the system and verify connections.
8. Run the operating-system-specific disk administration utility to add logical drive volumes. Some systems may require rebooting to use the new volumes.

For detailed information on drive migrations, go to the storage page on the following website: http://www.hp.com.

Adding a New Storage Enclosure

New storage enclosures with unconfigured drives are hot-pluggable. The following steps for attaching and installing a new enclosure to an existing MSA1000:

1. Install the new storage enclosure in the rack.

   Refer to the documentation shipped with the rack for instructions.

2. Attach SCSI cables to the MSA1000 expansion SCSI connectors and the storage enclosure.

   See Figure 42 for an illustration of the SCSI connectors.

3. Attach the power cords to the power supplies at the rear of the enclosure.
4. Configure new storage with the ACU or the CLI.

   Verify that additional storage is identified through the utility.
   Refer to the ACU or CLI chapters in this manual for further information.

5. Run the operating-system-specific disk administration utility to add logical drive volumes.
6. If necessary for your operating system, reboot server to use the new storage.

Updating the Controller Firmware

To load the latest version of the firmware for the MSA1000 Controller, go to the HP website at http://www.hp.com. From the HP StorageWorks MSA1000 page, select the Software, Firmware & Drivers option.

Note: All supported operating systems obtain MSA1000 Controller firmware upgrades from the HP website.

Installation instructions are available on the website.

You can use several methods to find out the version of the controller firmware you are currently running:

- In the CLI, use the `SHOW VERSION` command.
- On the MSA1000 Controller LCD panel, use the arrow keys to scroll backwards through the messages.

  Each time the MSA1000 storage system is restarted, the first message that is displayed includes the firmware version.
- In the ACU, highlight the controller and view the details.

Array Configuration Utility (ACU)

5

The Array Configuration Utility (ACU) is a browser-based utility that can be used any time the server is powered on. The ACU has different operating modes, allowing faster configuration or greater control over configuration options. The ACU suggests the optimum configuration for an unconfigured system and provides on-screen tips for individual steps of a configuration procedure. Online array capacity expansion, logical drive capacity extension, assignment of online spares, and RAID or stripe size migration are supported.

Note: The ACU is available for the following operating systems:

- Windows
- Linux
- NetWare (offline)

This chapter discusses:

- Installing the ACU on the Server, page 84
- Accessing the ACU, page 85
- Description of Screen Regions, page 88
- Configuring a New Controller, page 90
- Modifying an Existing Controller, page 100
- Probability of Logical Drive Failure, page 114

Note: The default host mode for servers accessing the MSA1000 is Windows. Servers running other operating systems must change the host mode to match their profile. See the "Selective Storage Presentation (SSP) " section for more information.

Installing the ACU on the Server

Windows and Linux environments can load the ACU onto their servers.

1. Download the ACU Smart Component from the HP website or from the MSA1000 Support Software CD supplied with the controller.
2. Install the ACU Smart Component onto the system.

   Installation instructions are included on the website and in the *hp StorageWorks Modular SAN Array 1000 Quick Installation Guide* that is included in the Modular SAN Array 1000 Setup and Management Kit included in the shipping carton of the MSA1000.

   When installation is complete, the ACU icon is displayed in the system tray.

Note: The minimum display settings for optimum performance are 1024 x 768 resolution and 256 colors. Refer to the *README.TXT* file for further information about browser and operating system support.

---

3. Click **Start** and navigate to **Programs > System Tools > Array Configuration Utility**.
4. Choose whether you would like to enable remote access.

Note: If remote access is disabled, the ACU can only be run on the server that has the Smart Component installed.

Accessing the ACU

You can run the ACU directly from the Smart Component or you can run it from Insight Manager 7.

Using the Smart Component

1. Open the browser.
   - If remote access is enabled, enter this text into the URL field (where *SERVERNAME* is the name or IP address of the host):

     `http://SERVERNAME:2301`
   - If remote access is disabled, enter this text into the URL field:

     `https://127.0.0.1:2381`

   The **System Management Home Page** opens.
2. Click **Storage Management** on the left side of the screen.

   The **ProLiant Storage Manager welcome screen** is displayed.

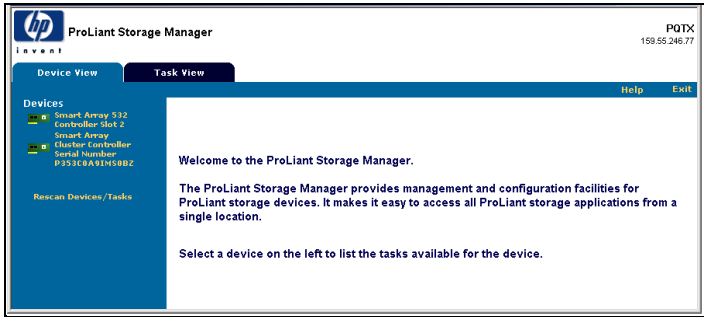

Figure 43: ProLiant Storage Manager welcome screen

3. Select a controller from the device list.
4. Click **Configure** in the main part of the screen.

   The ACU opens, and identifies the controllers that are connected to your system. This process may take a minute or two.

   When controller detection is complete, the configuration mode selection screen is displayed.

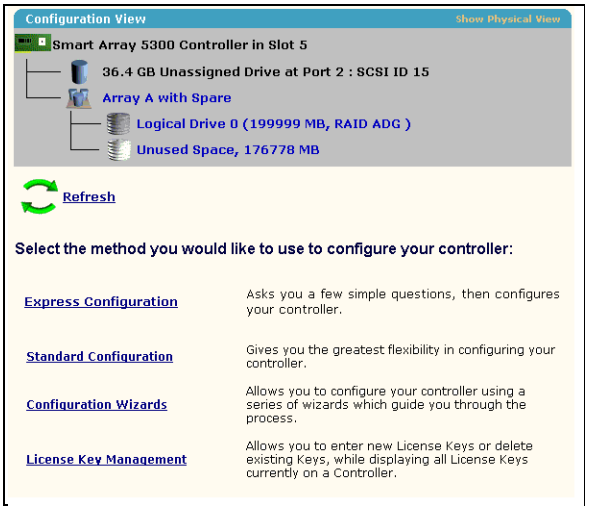
Configuration View
Show Physical View
Smart Array 5300 Controller in Slot 5
36.4 GB Unassigned Drive at Port 2 : SCSI ID 15
Array A with Spare
Logical Drive 0 (199999 MB, RAID ADG)
Unused Space, 176778 MB
Refresh
Select the method you would like to use to configure your controller:
Express Configuration
Asks you a few simple questions, then configures your controller.
Standard Configuration
Gives you the greatest flexibility in configuring your controller.
Configuration Wizards
Allows you to configure your controller using a series of wizards which guide you through the process.
License Key Management
Allows you to enter new License Keys or delete existing Keys, while displaying all License Keys currently on a Controller.

Figure 44: Configuration mode selection screen

The configuration mode selection screen provides up to three different methods that you can use to configure a controller. (**Express Configuration** mode is listed only if there is unused space or an unassigned drive on the selected controller.)

Note: If the controller that you selected is unconfigured, the gray Configuration View region of the screen shows only unassigned drives. Arrays, logical drives, and unused space are absent.

Details of the subsequent steps in the controller configuration process are given in the remainder of this chapter.

- If the controller is not configured (it has no arrays or logical drives, only unassigned physical drives), refer to the section, “Configuring a New Controller.”
- If the controller is already configured but you want to reconfigure it, refer to “Modifying an Existing Controller.”

Using Insight Manager 7

1. On the server where the ACU is located, confirm that the utility is configured to allow remote access.
2. On the remote system, connect to the Insight Manager server (port :280) and log in.
3. Select **Device Queries**.
4. Under **Device by Type**, select **All Servers**.
5. Connect to the server that is running the ACU.
6. Under **Device Links**, select the **System Management Home Page**.
7. Click **Storage Management** on the left side of the screen.

   The **ProLiant Storage Manager welcome screen** (Figure 43) is displayed.
8. Select a controller from the device list.
9. Click **Configure** in the main part of the screen.

   The ACU opens and identifies the controllers that are connected to your system. This process may take a minute or two.

   When controller detection is complete, the configuration mode selection screen (Figure 44) is displayed.

Description of Screen Regions

The appearance of a typical screen depends on which of the three configuration modes you use.

Typical Wizard-Based Screen

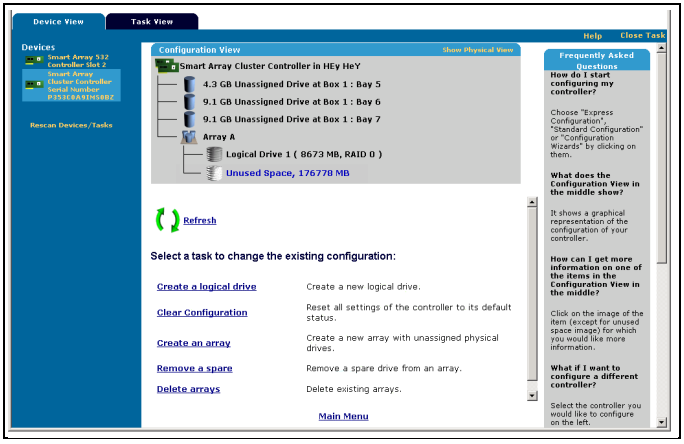
Device View
Task View
Help
Close Task
Devices
Smart Array 532 Controller Slot 2
Smart Array Cluster Controller Serial Number
Rescan Devices/Tasks
Configuration View
Show Physical View
Smart Array Cluster Controller in HEy HeY
4.3 GB Unassigned Drive at Box 1 : Bay 5
9.1 GB Unassigned Drive at Box 1 : Bay 6
9.1 GB Unassigned Drive at Box 1 : Bay 7
Array A
Logical Drive 1 (8673 MB, RAID 0)
Unused Space, 176778 MB
Refresh
Select a task to change the existing configuration:
Create a logical drive
Create a new logical drive.
Clear Configuration
Reset all settings of the controller to its default status.
Create an array
Create a new array with unassigned physical drives.
Remove a spare
Remove a spare drive from an array.
Delete arrays
Delete existing arrays.
Main Menu
Frequently Asked Questions
How do I start configuring my controller?
Choose "Express Configuration", "Standard Configuration" or "Configuration Wizards" by clicking on them.
What does the Configuration View in the middle show?
It shows a graphical representation of the configuration of your controller.
How can I get more information on one of the items in the Configuration View in the middle?
Click on the image of the item (except for unused space image) for which you would like more information.
What if I want to configure a different controller?
Select the controller you would like to configure on the left.

Figure 45: Typical wizard-based screen

The **Devices** list on the left side of the screen shows all the identifiable controllers that are connected to your system.

The gray **Configuration View** window in the upper central portion of the screen shows all arrays, logical drives, unused space, and unassigned physical drives that are connected to the selected controller. The logical view is shown by default.

- Click an icon to get further information about the corresponding item in a popup window (refer to Figure 46 for an example).
- Change to the physical configuration view at any time by clicking **Show Physical View** in the upper right corner of the window.

The **Main Menu** in the lower central portion of the screen shows the allowable options at this stage.

The **FAQ Column** on the right side of the screen lists information and tips that are relevant to the current screen. Check this region before clicking **Help** in the upper right-hand corner of the browser screen.

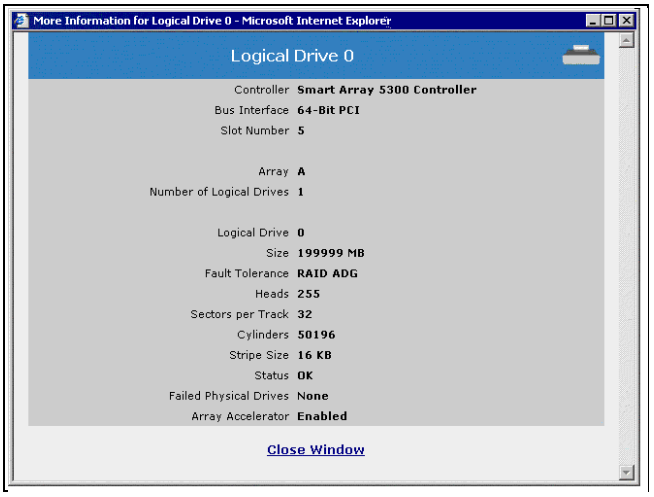

Figure 46: Typical More Information popup window

Typical Standard Mode Screen

This mode shows all the configuration options for a particular device at the same time, in a frame on the right side of the screen.

The FAQ column seen in the Wizard screens is absent.

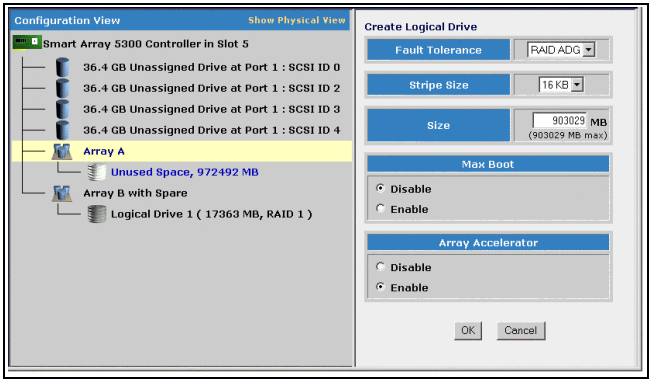

Figure 47: Typical standard mode screen

Configuring a New Controller

Open the ProLiant Storage Manager as described in "Accessing the ACU." Select a controller and click **Configure**. The configuration mode selection screen opens.

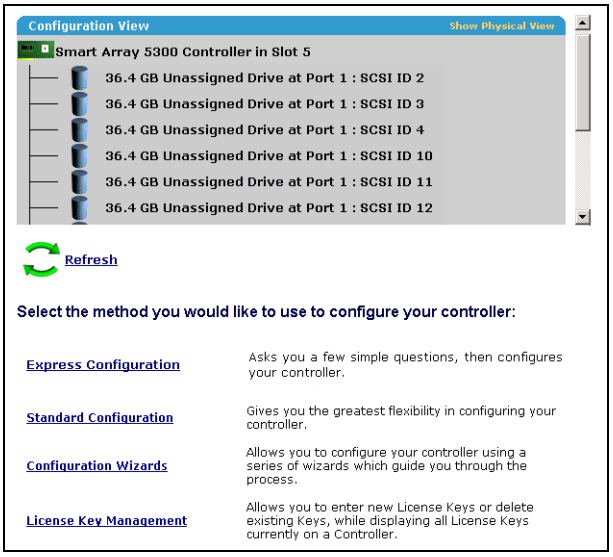
Configuration View
Show Physical View
Smart Array 5300 Controller in Slot 5
36.4 GB Unassigned Drive at Port 1 : SCSI ID 2
36.4 GB Unassigned Drive at Port 1 : SCSI ID 3
36.4 GB Unassigned Drive at Port 1 : SCSI ID 4
36.4 GB Unassigned Drive at Port 1 : SCSI ID 10
36.4 GB Unassigned Drive at Port 1 : SCSI ID 11
36.4 GB Unassigned Drive at Port 1 : SCSI ID 12
Refresh
Select the method you would like to use to configure your controller:
Express Configuration
Asks you a few simple questions, then configures your controller.
Standard Configuration
Gives you the greatest flexibility in configuring your controller.
Configuration Wizards
Allows you to configure your controller using a series of wizards which guide you through the process.
License Key Management
Allows you to enter new License Keys or delete existing Keys, while displaying all License Keys currently on a Controller.

Figure 48: Starting to configure a new controller

Note: If you select a controller that is already configured, arrays and logical drives are displayed in the **Configuration View** window in addition to any unassigned physical drives that may be present. The procedure for reconfiguring a previously configured controller is slightly more complex, and is described in "Modifying an Existing Configuration."

These methods are available for configuring the controller:

- **Express Configuration** automatically sets up the optimum configuration for the controller after you have answered a few simple questions.
- **Standard Configuration** lets you configure the controller manually without using wizards.
- **Configuration Wizards** guides you through a manual configuration process.

Using the Express Configuration Mode to Configure a New Controller

1. Click **Express Configuration**.

   The express mode start screen is displayed.

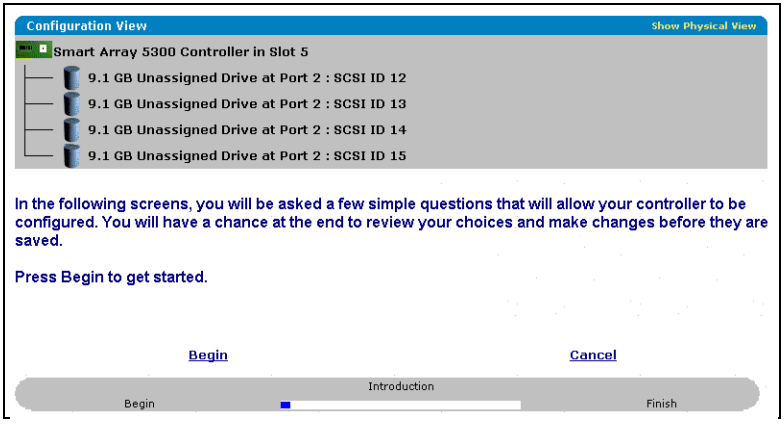
Configuration View
Show Physical View
Smart Array 5300 Controller in Slot 5
9.1 GB Unassigned Drive at Port 2 : SCSI ID 12
9.1 GB Unassigned Drive at Port 2 : SCSI ID 13
9.1 GB Unassigned Drive at Port 2 : SCSI ID 14
9.1 GB Unassigned Drive at Port 2 : SCSI ID 15
In the following screens, you will be asked a few simple questions that will allow your controller to be configured. You will have a chance at the end to review your choices and make changes before they are saved.
Press Begin to get started.
Begin
Cancel
Introduction
Begin
Finish

Figure 49: Express mode start screen

2. Click **Begin**.

   ACU creates the optimum number of arrays and logical drives from all of the physical drives that are attached to the controller.

   This process takes a few moments; when it is finished, the screen is updated. The gray **Configuration View** window shows the new configuration.

   Underneath this window is a list of possible fault tolerance levels for the first logical drive, as illustrated in Figure 50.

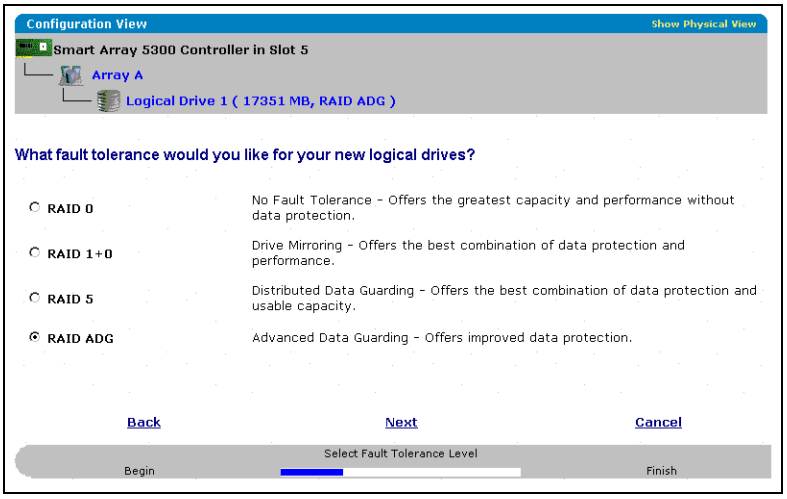
Configuration View
Show Physical View
Smart Array 5300 Controller in Slot 5
Array A
Logical Drive 1 (17351 MB, RAID ADG)
What fault tolerance would you like for your new logical drives?
RAID 0
No Fault Tolerance - Offers the greatest capacity and performance without data protection.
RAID 1+0
Drive Mirroring - Offers the best combination of data protection and performance.
RAID 5
Distributed Data Guarding - Offers the best combination of data protection and usable capacity.
RAID ADG
Advanced Data Guarding - Offers improved data protection.
Back
Next
Cancel
Select Fault Tolerance Level
Begin
Finish

Figure 50: Choosing a RAID level

3. Select a RAID level and click **Next**.

   If you have chosen a fault-tolerant RAID method, and an unassigned physical drive of the appropriate capacity is available, ACU asks if you want to assign a spare drive. Make your choice and click **Next**.

4. If ACU created more than one logical drive, you can now assign fault-tolerance levels and spare drives to the remaining logical drives, one logical drive at a time.

5. When you have assigned fault-tolerance levels and spare drives to all logical drives, the screen displays the chosen configuration and asks you to confirm that it is acceptable.

   - If you discard the suggested configuration, you are returned to the configuration mode selection screen so that you can configure the new array manually.
   - If you accept the suggested configuration, a screen is displayed confirming that ACU has saved the new configuration. At this point, you can configure another controller or you can exit ACU.

   Choose the appropriate radio button.

6. Click **Finish**.

Using the Configuration Wizards to Configure a New Controller

Create at least one array, and then populate the array with logical drives.

Creating an Array

1. Click **Configuration Wizards**.
2. Click **Create an array** and click **Begin**.

   The physical drive selection screen is displayed. (Use the scrollbars in the **Configuration View** region to see all of the physical drives and arrays.)

   A placeholder for the new array that you are about to create is already present.

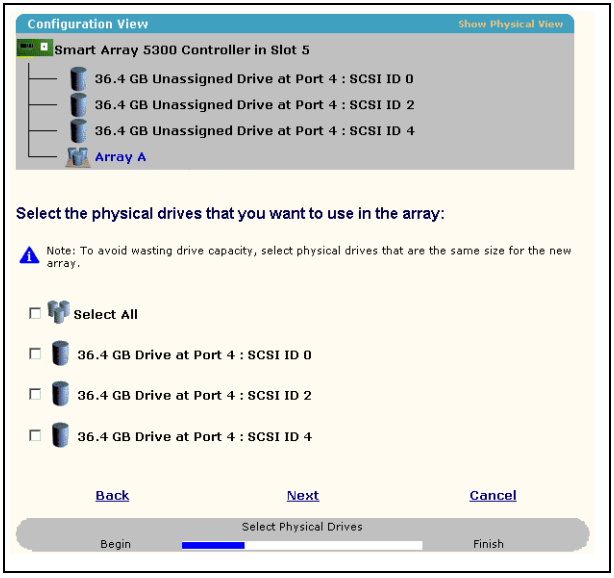
Configuration View
Show Physical View
Smart Array 5300 Controller in Slot 5
36.4 GB Unassigned Drive at Port 4 : SCSI ID 0
36.4 GB Unassigned Drive at Port 4 : SCSI ID 2
36.4 GB Unassigned Drive at Port 4 : SCSI ID 4
Array A
Select the physical drives that you want to use in the array:
Note: To avoid wasting drive capacity, select physical drives that are the same size for the new array.
Select All
36.4 GB Drive at Port 4 : SCSI ID 0
36.4 GB Drive at Port 4 : SCSI ID 2
36.4 GB Drive at Port 4 : SCSI ID 4
Back
Next
Cancel
Select Physical Drives
Begin
Finish

Figure 51: Physical drive selection screen

3. Choose the physical drives that you want to use in the array.
   - Use physical drives of comparable capacity.

     ACU uses the same amount of space from each physical drive to build an array. Because this amount cannot exceed the capacity of the smallest drive, the extra capacity of any larger drive in the array is unused.
   - For better system performance, use physical drives that are connected to different ports on the controller.

- In RAID 5 configurations, keep the risk of logical drive failure low by assigning no more than 14 physical drives to the array.

Each time you add a physical drive to the array, the configuration view is updated to show how much free space remains on the array.

4. Click **Next** when you have finished adding physical drives to the array.
5. If an unassigned physical drive of the appropriate capacity is available, the ACU asks you whether you want to assign spare drives to the array.
   - If you do not want this array to have a spare, click **No** and click **Next**.
   - To assign spare drives to the array, click **Yes** and click **Next**.

     On the next screen, select the drives that you want to be the spares and click **Next**.

Note: Assigning several spares to an array lets you postpone replacement of faulty drives but it does not increase the fault-tolerance level of any logical drives in the array. For example, a logical drive in a RAID5 configuration suffers irretrievable data loss if two physical drives fail, regardless of the number of spare drives assigned.

Note: A spare drive can be shared by several arrays.

---

6. Click **Finish** to confirm the configuration.

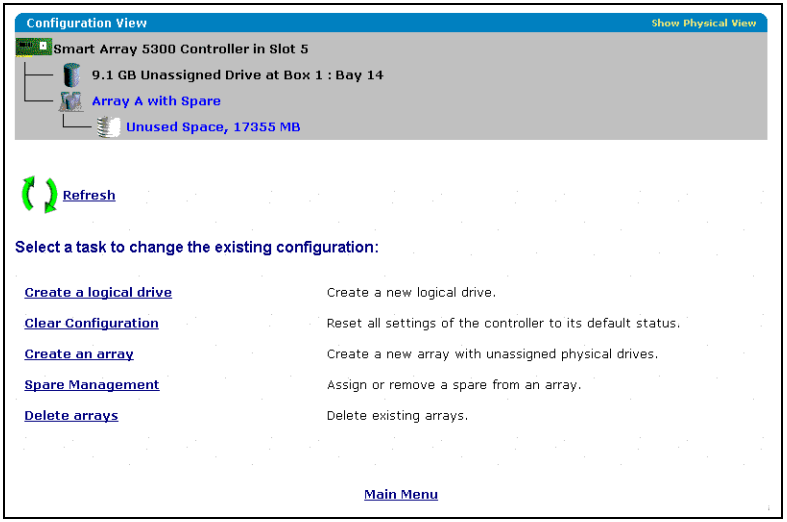
Configuration View
Show Physical View
Smart Array 5300 Controller in Slot 5
9.1 GB Unassigned Drive at Box 1 : Bay 14
Array A with Spare
Unused Space, 17355 MB
Refresh
Select a task to change the existing configuration:
Create a logical drive
Create a new logical drive.
Clear Configuration
Reset all settings of the controller to its default status.
Create an array
Create a new array with unassigned physical drives.
Spare Management
Assign or remove a spare from an array.
Delete arrays
Delete existing arrays.
Main Menu

Figure 52: Configured array screen (no logical drives)

7. To create more arrays on the same controller, repeat the previous steps.

Creating a Logical Drive

1. Click **Create a logical drive** and click **Begin**.
2. Select an array that has unused space and click **Next**.

   (A logical drive can only be configured on an array that has unused drive space available.)
3. Select the fault tolerance level for the logical drive and click **Next**.

   Only RAID levels that are possible for this configuration are listed. For example, RAID 5 is not listed if the array has only two physical drives.
4. Select the stripe size and click **Next**.

   The default stripe size gives optimum performance in a mixed read/write environment. If your system is used in a different environment, use the following table to determine what stripe size to set.

Table 12: Optimum Stripe Size

Type of Server Application	Suggested Strip Size Change
Mixed read/write	Accept the default value
Mainly sequential read (such as audio/video applications)	Use larger stripe sizes for best performance
Mainly write (such as image manipulation applications)	Use smaller stripes for RAID 5, RAID ADG* Use larger stripes for RAID 0, RAID 1+0
*Not all controllers support RAID ADG	

5. Decide whether to use MaxBoot and click **Next**.

   MaxBoot increases the number of sectors used per track from 32 to 63. This increase allows a larger boot partition for operating systems that use cylinders, heads, and sectors of a physical drive to determine the drive size. MaxBoot also lets you create a larger logical drive or increase the logical drive size (extend it) at some later time.

   Logical drive performance is likely to decrease with MaxBoot enabled.
6. Set the size that you want the logical drive to be and click **Next**.

   The default size shown is the largest possible logical drive size for the RAID level that you chose and the set of physical drives that is being used. Reducing the size of the logical drive liberates drive space, which you can use to build additional logical drives on the same array.

7. If the controller has an array accelerator, a screen is now displayed that lets you disable it for the currently selected logical drive.

   Choose whether to disable the array accelerator and click **Next**.

Note: Disabling the array accelerator for a logical drive reserves use of the accelerator cache for other logical drives on the array that need to have the maximum possible performance (such as those that contain database information).

   The gray **Configuration View** window shows the configuration that you have chosen.

8. Check that the configuration is acceptable and click **Finish**.

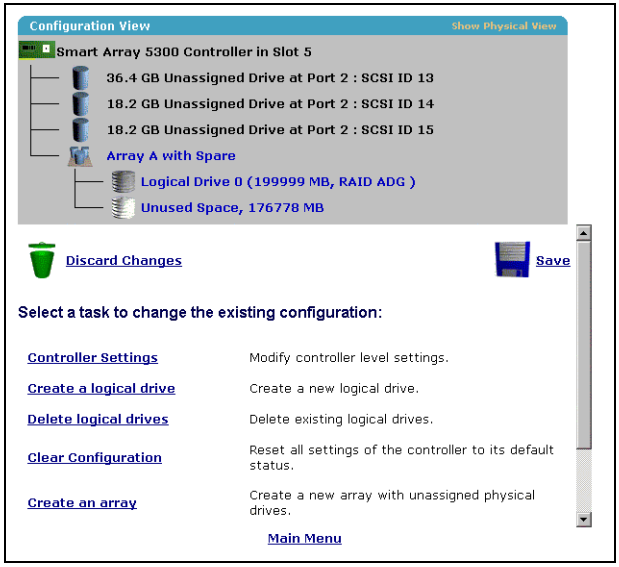
Configuration View
Show Physical View
Smart Array 5300 Controller in Slot 5
36.4 GB Unassigned Drive at Port 2 : SCSI ID 13
18.2 GB Unassigned Drive at Port 2 : SCSI ID 14
18.2 GB Unassigned Drive at Port 2 : SCSI ID 15
Array A with Spare
Logical Drive 0 (199999 MB, RAID ADG)
Unused Space, 176778 MB
Discard Changes
Save
Select a task to change the existing configuration:
Controller Settings
Modify controller level settings.
Create a logical drive
Create a new logical drive.
Delete logical drives
Delete existing logical drives.
Clear Configuration
Reset all settings of the controller to its default status.
Create an array
Create a new array with unassigned physical drives.
Main Menu

Figure 53: New logical drive before saving

9. Click the **Save** icon to commit the changes to the controller and click **OK** on the confirmation alert. (If you discard the changes, all changes since the previous save are lost.)

To change the array configuration, refer to "Modifying an Existing Controller."

Using the Standard Configuration Mode to Configure a New Controller

1. In the configuration mode selection screen, click **Standard Configuration**.
2. Click an item in the **Configuration View** window.

   The screen displays a list of the tasks that are available for that item.

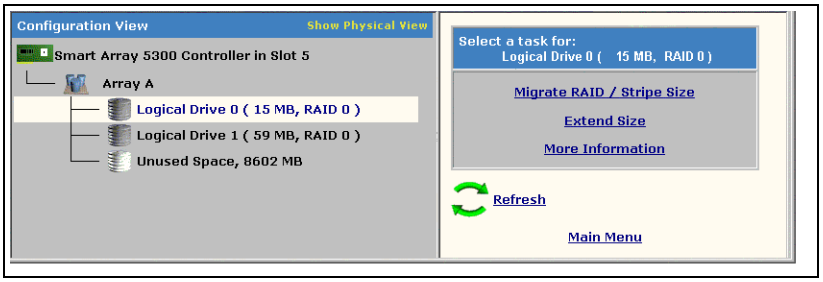
Configuration View
Show Physical View
Smart Array 5300 Controller in Slot 5
Array A
Logical Drive 0 (15 MB, RAID 0)
Logical Drive 1 (59 MB, RAID 0)
Unused Space, 8602 MB
Select a task for:
Logical Drive 0 (15 MB, RAID 0)
Migrate RAID / Stripe Size
Extend Size
More Information
Refresh
Main Menu

Figure 54: Typical task list for a logical drive

   The listed (available) tasks are a subset of the total number of tasks that are possible for the selected item. Which of the possible tasks are listed for an item and which are omitted depends on the controller model and configuration. (For example, if the selected controller has no unassigned physical drives, **Create Array** is not an available task.)

   Table 13 lists all the possible tasks for every type of item.

3. Click a task link.

   A list of all possible configuration options for that task is displayed on the right-hand side of the screen (replacing the task list). See the “Description of Screen Regions” section for an example of this type of screen.

4. Set the configuration options.
5. Click **OK**.

Table 13: Possible Tasks in Standard Configuration Mode

Item	Tasks
Controller	Clear Configuration Controller Settings Create Array Logical Drive Array Accelerator Settings Selective Storage Presentation (for MSA1000 controllers) More Information
Array	Assign Spare Create Logical Drive Delete Expand Remove Spare More Information
Logical Drive	Delete Extend Size Migrate RAID/Stripe Size Selective Storage Presentation (for RA4x00 controllers)
Unused space	There are no available tasks associated with this item.
Note: The **More Information** task is present for all items except unused space. When you click this task link, a pop-up window is displayed with additional information about the selected item.	

Modifying an Existing Controller

Open the ProLiant Storage Manager as described in "Accessing the ACU," select the controller you want to reconfigure, and click **Configure**.

The configuration mode selection screen opens.

Depending on the existing configuration of the controller, this screen may give you the choice of using Express mode, Standard mode, or the Wizards mode to reconfigure arrays.

When you select an MSA1000 Controller, another link on this screen is displayed to allow you to configure switches.

Using the Express Configuration Mode to Modify an Existing Controller

This mode is available only if the selected controller has unused space on an array or unassigned physical drives.

1. Click **Express Configuration** and click **Begin**.

   If there are unassigned physical drives on the controller, you can create a new array or expand an existing array. Make your choice and click **Next**.

Note: The expansion process takes about 15 minutes per gigabyte, or considerably longer if the controller does not have a battery-backed cache. While array expansion is occurring, no other expansion, extension, or migration can occur simultaneously on the same controller.

   The screen displays the optimum configuration for the controller and asks you to confirm that it is acceptable.

2. Choose the appropriate radio button and click **Finish**.

Using the Configuration Wizards to Modify an Existing Controller

The options listed in the menu region of the screen depend on the controller model and configuration. For example, the **Expand array** option is listed only if there is at least one unassigned physical drive connected to the controller.

These are the possible menu options:

- Clear Configuration
- Controller Settings
- Create an array
- Create a logical drive
- Delete arrays
- Delete logical drives
- Expand array
- Extend logical drive
- Migrate a logical drive
- Spare Management
- Selective Storage Presentation

Clear Configuration

This task deletes all logical drives connected to the controller, reconfigures the arrays into independent (unassigned) physical drives, and resets all controller settings to their default values.

1. Click **Clear Configuration** and click **Begin**.

   A warning screen is displayed to remind you that you will lose all data on the logical drive.
2. Click **Delete** to continue.
3. Click **Finish** to accept the changes.
4. Click **Save** to apply the changes to the system and click **OK** on the confirmation alert.

The physical drives are now available for reconfiguration.

Controller Settings

The default controller settings provided by ACU are adequate for many purposes. However, this task lets you alter the priority settings for array expansion and rebuild. You can also disable the array accelerator (if one is present) or change the ratio of read cache to write cache (if the controller has battery-backed cache).

To change the controller settings:

1. Click **Controller Settings** and click **Begin**.

   The next two screens let you change the expand priority and the rebuild priority settings. These settings determine how much importance you want an array expansion or rebuild to have relative to normal I/O operations.

   - With **low** priority, the expansion or rebuild takes place only when the array controller is not busy handling normal I/O requests. This setting has minimal effect on normal I/O operations. However, there is an increased risk that data will be lost if another physical drive fails while the rebuild is in progress.
   - With **high** priority, the rebuild or expansion occurs at the expense of normal I/O operations. Although system performance is affected, this setting provides better data protection because the array is vulnerable to additional drive failures for a shorter time.
   - At the **medium** priority setting, expansion or rebuild occurs for half of the time, and normal I/O requests are handled during the rest of the time.

2. Set the expand priority to high, medium, or low; click **Next**.
3. Set the rebuild priority and click **Next**.
4. If the controller has an array accelerator, a screen is now displayed that lets you disable it for particular logical drives.

   Choose whether to disable the array accelerator for any logical drives and click **Next**.

Note: Disabling the array accelerator for a particular logical drive reserves use of the accelerator cache for other logical drives on the array that need to have maximum possible performance (such as those that contain database information).

---

5. If the controller has a battery-backed cache, a screen is now displayed that lets you change the read/write cache ratio.

   Choose the ratio that you want the controller to use and click **Next**.

Note: This ratio determines the amount of memory allocated to read and write operations. Different types of applications have different optimum ratios. You can only change the ratio if the controller has a battery-backed cache (only battery-backed cache can be used for write cache) and if there are logical drives configured on the controller.

6. Click **Finish** to accept the changes.
7. Click **Save** to apply the changes and click **OK** at the confirmation alert.

Create an Array

1. Click **Create an array** and click **Begin**.
2. Choose the physical drives you want to use in the array.
   - Use physical drives of comparable capacity.

     ACU uses the same amount of space from each physical drive to build an array. Because this amount cannot exceed the capacity of the smallest physical drive, the extra capacity of any larger drive in the array is unused.
   - For better system performance, use physical drives that are attached to different ports on the controller.
   - In RAID 5 configurations, keep the risk of logical drive failure low by assigning no more than 14 physical drives to the array.

   Each time that you add a physical drive to the array, the configuration view is updated to show how much free space remains on the array.
3. Click **Next** when you have finished adding physical drives to the array.

   If a spare or unassigned physical drive of the appropriate capacity is available, ACU asks you whether you want to assign a spare drive to the array.
   - If you do not want this array to have a spare, click **No** and click **Next**.
   - To assign spare drives to the array, click **Yes** and click **Next**. On the next screen, select the drives that you want to be the spares and click **Next**.

Note: A spare can be shared by several arrays.

4. Click through the remaining screens to confirm the configuration.

Create a Logical Drive

1. Click **Create a logical drive** and click **Begin**.
2. Choose an array that has unused space and click **Next**.
3. Choose the fault-tolerance level that you want the logical drive to use and click **Next**.

   Only RAID levels that are possible for this configuration are listed. For example, RAID 5 is not listed if the array has only two physical drives.
4. Choose the stripe size and click **Next**.

   The default stripe size gives optimum performance in a mixed read/write environment. If your system is used in a different environment, see Table 14 to determine what stripe size to set.

Table 14: Optimum Stripe Size

Type of Server Application	Suggested Stripe Size Change
Mixed read/write	Accept the default value
Mainly sequential read (such as audio/video applications)	Use larger stripe sizes for best performance
Mainly write (such as image manipulation applications)	Use smaller stripes for RAID 5, RAID ADG* Use larger stripes for RAID 0, RAID 1+0
*Not all controllers support RAID ADG	

5. Decide whether to use MaxBoot and click **Next**.

   MaxBoot increases the number of sectors used per track from 32 to 63. This increase allows a larger boot partition for operating systems that use cylinders, heads, and sectors of a physical drive to determine the drive size. MaxBoot also lets you create a larger logical drive or increase the logical drive size (extend it) at some later time.

   Logical drive performance is likely to decrease with MaxBoot enabled.
6. Set the size that you want the logical drive to be and click **Next**.

   The default size shown is the largest possible logical drive size for the RAID level that you chose and the set of physical drives that is being used. Reducing the size of the logical drive liberates drive space, which you can use to build additional logical drives on the same array.

7. If the controller has an array accelerator, a screen is now displayed that lets you disable it for the currently selected logical drive.

   Make your choice and click **Next**.

Note: Disabling the array accelerator for a particular logical drive reserves use of the accelerator cache for other logical drives on the array that need the maximum possible performance (such as those that contain database information.)

   The gray **Configuration View** window shows the configuration that you have chosen.

8. Check that the configuration is acceptable and click **Finish**.
9. Click **Save** to apply the changes to the system and click **OK** on the confirmation alert.

Delete Arrays

This task deletes logical drives on an array and converts the array into a group of unassigned physical drives. You can then reconfigure the unassigned physical drives into one or more new arrays, or you can use the liberated physical drive space for expansion of another array on the same controller.

1. Click **Delete arrays** and click **Begin**.
2. Select the arrays that you want to delete and click **Next**.

   A warning screen is displayed to remind you that you will lose all data on the array.

3. Click **Delete** to continue and click **Finish** to accept the changes.
4. Click **Save** to apply the changes to the system and click **OK** on the confirmation alert.

Delete Logical Drives

This task deletes the selected logical drive and converts it into unused drive space. You can then use this unused drive space to:

- Create new logical drives.
- Migrate the RAID level or stripe size of an existing logical drive.
- Extend an existing logical drive on the same array, if the operating system allows logical drive extension.

To delete a logical drive:

1. Click **Delete logical drives** and click **Begin**.
2. Select the logical drives that you want to delete and click **Next**.

   A warning screen is displayed to remind you that you will lose all data on the logical drive.
3. Click **Delete** to continue and click **Finish** to accept the changes.
4. Click **Save** to apply the changes to the system and click **OK** on the confirmation alert.

Expand Array

Note: **Expand array** is listed only if there is an unassigned physical drive on the controller. The unassigned drive must also have a capacity no less than that of a drive in an existing array. If these conditions are not fulfilled, install at least one suitable physical drive on the controller, and then click **Refresh.**

This task increases the storage capacity of an existing array. You can use the additional storage space to:

- Create new logical drives.
- Migrate the RAID level or stripe size of existing logical drives.
- Extend existing logical drives on the array, if the operating system allows logical drive extension.

Note: The expansion process takes about 15 minutes per gigabyte, or considerably longer if the controller does not have a battery-backed cache. While array expansion is taking place, no other expansion, extension, or migration can occur simultaneously on the same controller.

1. Click **Controller Settings** and check that the expand priority setting is acceptable.
2. Back up all data on the array.

   Although array expansion is unlikely to cause data loss, observing this precaution provides additional data protection.
3. Click **Expand array** and click **Begin**.
4. Choose the array that you want to expand and click **Next**.
5. Select the physical drives that you want to add to the array and click **Next**.
6. Click **Finish** to accept the changes.

   At this point (before clicking **Save** in the next step), you can create logical drives on the unused space created by the expansion. You can also arrange to expand another array on the same controller by repeating the previous steps. However, the controller can expand only one array at a time; remaining array expansions are queued.
7. Click **Save**.

   The controller now rearranges (re-stripes) the existing logical drives and their data so that they extend over all the physical drives in the enlarged array.

To check the progress of an array expansion, click the icon for that array in the **Configuration View** window. A **More Information** popup window is displayed that lists the drive status.

Extend Logical Drive

This option increases the storage capacity of a logical drive by adding unused space on an array to a logical drive on the same array. The unused space is obtained either by array expansion (refer to the Expand Array section), or by deleting another logical drive on the same array.

Not all operating systems support online logical drive extension through ACU. Also, **offline** logical drive extension is possible for some operating systems by backing up data, reconfiguring the array, and restoring data from backup. Check the operating system documentation for current information.

Note: The extension process takes about 15 minutes per gigabyte, or considerably longer if the controller does not have a battery-backed cache. While logical drive extension is taking place, no other expansion, extension, or migration can occur simultaneously on the same controller.

---

1. Back up all data on the logical drive.

   Although logical drive extension is unlikely to cause data loss, observing this precaution provides additional data protection.

2. Click **Extend logical drive** and click **Begin**.
3. Select the logical drive that you want to extend and click **Next**.
4. Type the new size of the logical drive into the size field.
5. Click **Finish.**

   At this point (before clicking **Save** in the next step), you can arrange to extend another logical drive on the same controller by repeating the previous steps. However, the controller can extend only one logical drive at a time; remaining extensions are queued.

6. Click **Save**.

   Logical drive extension begins.

To check the progress of a logical drive extension, click the icon for that logical drive in the **Configuration View** window. A **More Information** popup window is displayed that describes the drive status.

Migrate a Logical Drive

This option lets you alter the stripe size (data block size) or RAID level, or both, for a selected logical drive. There might need to be unused drive space available on the array for the migration to be possible, depending on the initial and final settings for the stripe size and RAID level.

Note: The migration process takes about 15 minutes per gigabyte, or considerably longer if the controller does not have a battery-backed cache. While migration is taking place, no other expansion, extension, or migration can occur simultaneously on the same controller.

---

1. Back up all data on the logical drive.

   Although migration is unlikely to cause data loss, observing this precaution provides additional data protection.
2. Click **Migrate a logical drive** and click **Begin**.
3. Select the logical drive and click **Next**.
4. Select the new RAID level and click **Next**.

   Only RAID levels that are possible for this configuration are shown. For example, RAID 5 is not listed if the array has only two physical drives.
5. Select the stripe size and click **Finish** to accept the changes. (Only stripe sizes that are possible for this configuration are shown.)

   At this point (before clicking **Save** in the next step), you can arrange to migrate another logical drive on the same controller by repeating the previous steps. However, the controller can migrate only one logical drive at a time; remaining migrations are queued.
6. Click **Save**.

   Migration begins.

To check the progress of a migration, click the icon for that logical drive in the **Configuration View** window. A **More Information** popup window is displayed that lists the drive status.

Spare Management

Note: An array can have several spares, and a spare can be shared by several arrays.

1. Click **Spare Management** and click **Begin**.
2. Select the array that needs to have spare drives added or removed.
3. Select the drives that you want to assign as spares, and deselect the appropriate checkboxes for spares that you want to remove.

Note: Assigning several spares to an array lets you postpone replacement of faulty drives but it does not increase the fault-tolerance level of any logical drives in the array. For example, a logical drive in a RAID5 configuration suffers irretrievable data loss if two physical drives fail, regardless of the number of spare drives assigned.

4. Click **Next**.
5. Click **Finish** to accept the changes.
6. Click **Save** to apply the changes to the system and click **OK** on the confirmation alert.

Selective Storage Presentation (SSP)

This menu option lets you forbid selected host controllers from accessing a logical drive. This prevents data corruption that may occur when servers using different operating systems access the same data.

1. Click **Selective Storage Presentation** and click **Begin**.

   A screen is displayed that lets you enable or disable SSP.

   - Disabling SSP lets all host controllers gain access to the selected logical drive.
   - Enabling SSP lets you deny access to selected hosts.

2. Select the appropriate radio button and click **Next**.

   If you select **Enable**, the screen lists all identified host controllers. Select the host controllers that are to have access to each logical drive, define the host mode for each controller, rename the connections if necessary, and click **Next**.

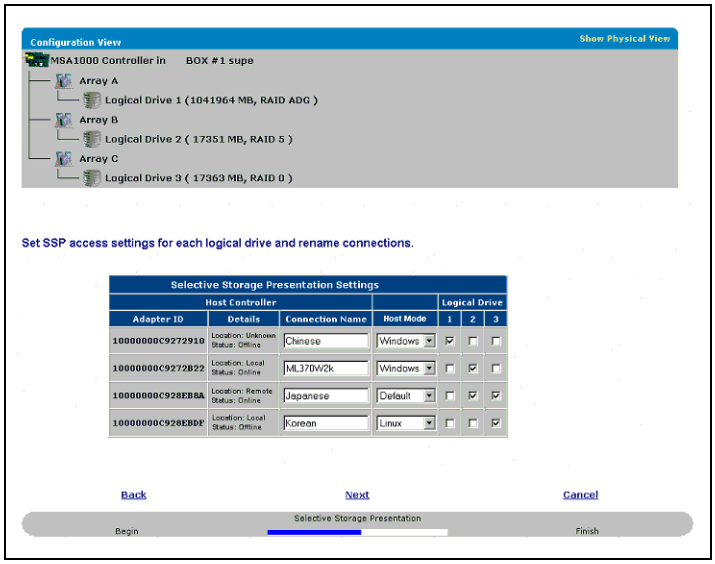

Figure 55: Typical SSP screen for MSA 1000

3. Click **Finish**.

Using Standard Configuration Mode to Modify an Existing Controller

1. Click **Standard Configuration**.
2. Click an item in the **Configuration View** window.

   The right side of the screen displays a list of the tasks that are available for that item.

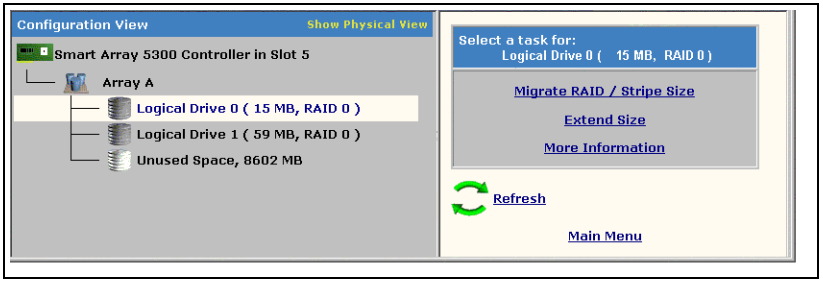
Configuration View
Show Physical View
Smart Array 5300 Controller in Slot 5
Array A
Logical Drive 0 (15 MB, RAID 0)
Logical Drive 1 (59 MB, RAID 0)
Unused Space, 8602 MB
Select a task for:
Logical Drive 0 (15 MB, RAID 0)
Migrate RAID / Stripe Size
Extend Size
More Information
Refresh
Main Menu

Figure 56: Typical task list for a logical drive

   The listed (available) tasks are a subset of the total number of tasks that are possible for the selected item. Which of the possible tasks are listed for an item and which are omitted depends on the controller model and configuration. (For example, if the selected controller has no unassigned physical drives, **Create Array** is not an available task.)

   Table 15 lists the possible tasks for every type of item.

3. Click a task link.

   A list of all possible configuration options for that task is displayed on the right-hand side of the screen (replacing the task list). Refer to the "Description of Screen Regions" section for an example of this type of screen.

4. Set the configuration options.
5. Click **OK**.

Table 15: Possible Tasks in Standard Configuration Mode

Item	Available Task
Controller	Clear Configuration Controller Settings Create Array Logical Drive Array Accelerator Settings Selective Storage Presentation (for MSA1000 and Smart Array Cluster Storage controllers) More Information
Array	Assign Spare Create Logical Drive Delete Expand Remove Spare More Information
Logical Drive	Delete Extend Size Migrate RAID / Stripe Size Selective Storage Presentation (for RA4x00 controllers) More Information
Unused space	No available tasks are associated with this item.
Note: The **More Information** task is present for all items except unused space. When you click this task link, a popup window is displayed with additional information about the selected item.	

Probability of Logical Drive Failure

The probability that a logical drive will fail depends on the RAID level setting. If the logical drive does not have an online spare, the following results apply.

- A RAID 0 logical drive fails if only one physical drive fails.
- A RAID 1+0 logical drive fails if any two failed physical drives are mirrored to each other.
  - The **maximum** number of physical drives that can fail **without** causing failure of the logical drive is *n*/2, where *n* is the number of hard drives in the array. In practice, a logical drive usually fails before this maximum is reached. As the number of failed drives increases, it becomes increasingly unlikely that a newly failed drive is not mirrored to a previously failed drive.
  - The **minimum** number of physical drive failures that can cause the logical drive to fail is two, **if** the two drives happen to be mirrored to each other. The probability that this will happen decreases as the total number of hard drives in the array increases.
- A RAID 5 logical drive fails if two physical drives fail.
- A RAID ADG logical drive fails when three physical drives fail.

At any given RAID level, the probability of logical drive failure increases as the number of physical drives in the logical drive increases.

The graph in Figure 57 provides more quantitative information. The data for this graph is calculated from the mean time between failure (MTBF) value for a typical physical drive, assuming that no online spares are present. If an online spare is added to any of the fault-tolerant RAID configurations, the probability of logical drive failure is further decreased.

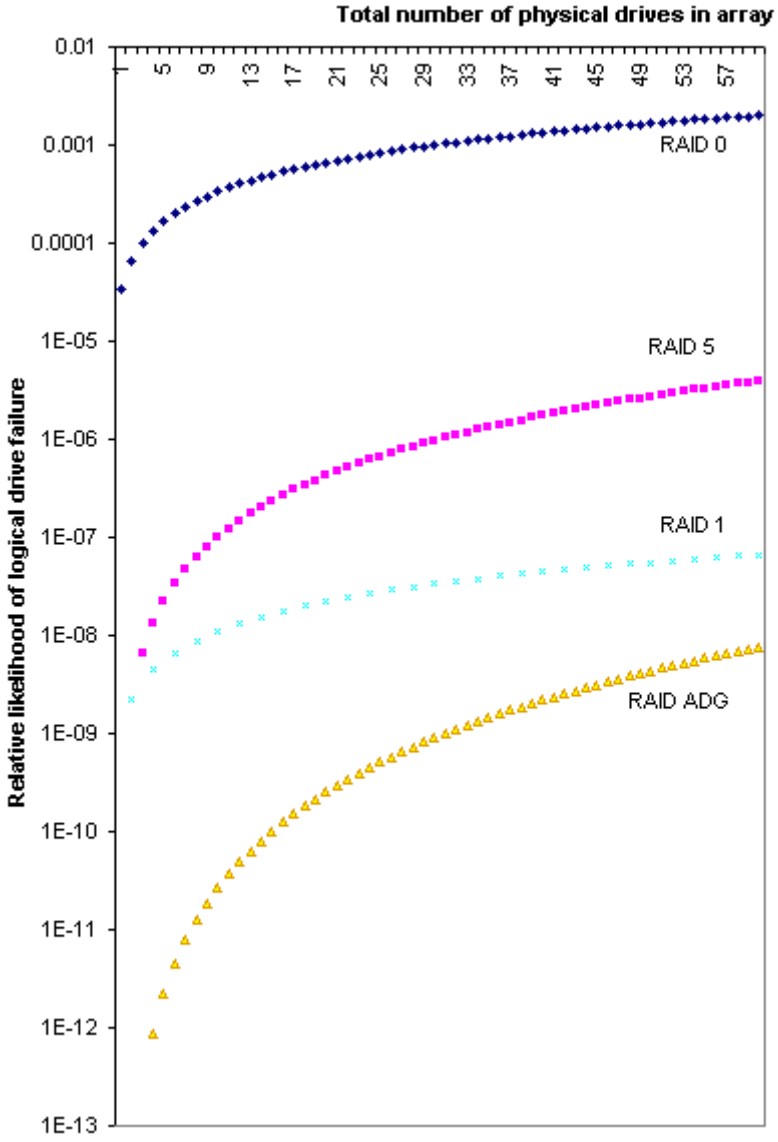

Figure 57: Relative probability of a logical drive failure

Command Line Interface (CLI)

6

The Command Line Interface (CLI) is used to configure and manage the MSA1000. In addition, the CLI can be used to display system setup information and status. It may also provide information on devices that are attached to the controller.

Note: The CLI is available for all supported operating systems.

The CLI is accessed through a host computer connected to the serial port of the MSA1000.

Some CLI configuration and management tasks include configuring storage units (LUNs), setting the addressing mode, limiting access to storage, and viewing information on MSA1000 components (controller, unit, and cache).

This chapter discusses:

- CLI Overview, page 118
- CLI Setup, page 121
- Help Commands, page 124
- Display Commands, page 127
- Array Controller Configuration Commands, page 134
- LUN Management Commands, page 137
- Server Connection Commands, page 152
- Selective Storage Presentation/Access Control List Commands, page 156

CLI Overview

After accessing the interface, enter a command string at the CLI prompt (`CLI>`). Commands must use a specific, preset syntax. After a declarative command is entered and executed, the results are displayed at the CLI prompt.

When using the CLI, the following keystrokes have special meaning:

Table 16: CLI Special Keys

Keyboard Keys	Arrow Key	Meaning
Ctrl_B	left arrow	move the cursor back one character
Ctrl_F	right arrow	move the cursor forward one character
Ctrl_P	up arrow	recall the previous command in the command buffer
Ctrl_N	down arrow	recall the next command in the command buffer

Note: The CLI uses a zero-based numbering system. For example, LUN number assignments begin with 0.

When the MSA1000 has two controllers, the same firmware image is run on both controllers. The two controllers communicate with each other through a PCI bus called the inter-controller link (ICL). Each controller has a serial port and has a CLI available to the users who connect the serial port to a serial terminal.

Some CLI commands are entered from one controller's CLI prompt but are for the other controller to execute. Command syntax instructs the controller to accept user input, pass the command to the other controller, and then display the result.

The following keywords are used in the CLI to indicate a specific controller:

- `this_controller`—is included in a command syntax to refer to the controller that the CLI is connected to.
- `other_controller`—is included in a command syntax to refer to the other controller in the MSA1000.

CLI Command Syntax

As previously mentioned, CLI commands are entered at the CLI prompt. Commands are **not** case sensitive and must be typed out in full.

CLI command strings include the basic command plus specific command options, some of which are mandatory and some of which are optional.

The CLI does not support line-continuation characters. If all characters of a command do not fit on one line of the CLI, let them wrap around to the next line on the screen. The maximum command length is 255 characters.

Example Command

```
ADD UNIT 0 DATA="DISK101-DISK103" RAID_LEVEL=0
```

This example command has three portions—the basic command, plus two command options.

Basic Command

```
add unit
```

The basic command includes a word or phrase used to instruct the controller. Commands usually contain a verb with a noun. Every CLI command must begin with a basic command.

Command Options

```
0
data="disk101-disk103"
raid_level=0
```

An option is defined as words or phrases listed after the basic command that supply necessary information to support the command.

Some CLI commands require these parameters. If parameters are required but are not entered, the CLI command string is considered invalid. The syntax of the parameter value is uniquely defined for each CLI command, but must be less than 20 alphanumeric characters.

Some CLI commands support options that modify the command, but are not required. If a modifying option is available but not used, a default value will be used.

Overview of CLI Storage Configuration Procedures

When using the CLI to initially configure the MSA1000 Controller and its storage, use the following sequence:

1. Enter array controller settings, including global parameters.

   See "Array Controller Configuration Commands" for command descriptions.

Note: OpenVMS environments must assign a unique ID number to each controller in the MSA1000. See "Setting Global Variables" for instructions.

---

2. Create the LUNs.

   See "LUN Management Commands" for command descriptions.

Note: OpenVMS environments must assign a unique ID number to each LUN. See "Assigning a Name or ID to a LUN" for instructions.

---

3. Enter server connection information to indicate the HBAs that will access the MSA1000.

   See "Server Connection Commands" for command descriptions.

Note: Each server accessing the MSA1000 must identify its profile type. See "Managing Connection Profiles" for instructions.

---

4. (Optional) Enter Selective Storage Presentation (SSP) and Access Control List (ACL) information.

   See "Selective Storage Presentation/Access Control List Commands" for command descriptions.

CLI Setup

Before using the serial interface to connect to the MSA1000, a serial connection to the controller must be established.

Note: Any terminal emulator program may be used, but the following instructions demonstrate setting up the serial connection with HyperTerminal.

Setting up the Serial Connection

To access the CLI on the MSA1000:

1. Make sure there is an MSA1000 Controller installed in Slot 1 of the MSA1000. (Slot 1 is located on the front right of the MSA1000.)
2. Connect the MSA1000 to a host computer using the custom serial cable included in the shipping carton of the MSA1000.

   An additional or replacement custom serial cable can be ordered using part number 259992-001.

   - MSA1000 adapter: RJ-45Z
   - Host computer adapter: serial port on the host computer
3. Set up a terminal emulator.

   To set up HyperTerminal:

   a. Access HyperTerminal.

      If you need to load HyperTerminal on the server, navigate the Web and download the latest version.

      If HyperTerminal is already loaded on the server, but its location is unknown, to locate this program on a Windows-based machine, go to **Start**, **Search**, **For Files and Folders**. Enter `hypertrm.exe` as search term.

      After this program is installed or located, if desired, create a shortcut for it to simplify subsequent accesses.

   b. Open HyperTerminal.

      When accessed for the first time, the **New Connection** dialog box is displayed.

      See Figure 58 for an example of the New Connection dialog box.

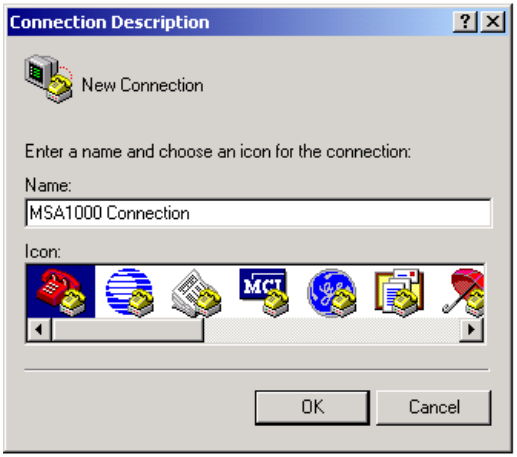

Figure 58: HyperTerminal New Connection dialog box

c. In the **New Connection** dialog box, type a name to associate with the connection between the MSA1000 and the host computer. Then click **OK**.

The **Connect To** dialog box is displayed.

See Figure 59 for an example of the **Connect To** dialog box.

Figure 59: HyperTerminal Connect To dialog box

d. In the **Connect To** dialog box, expand the **Connect using** drop-down box, select the appropriate COM port, and click **OK**.

   Figure 59 illustrates selecting **COM1**.

   The **COM Properties** dialog box is displayed.

   Figure 60 is an example of the COM1 Properties dialog box.

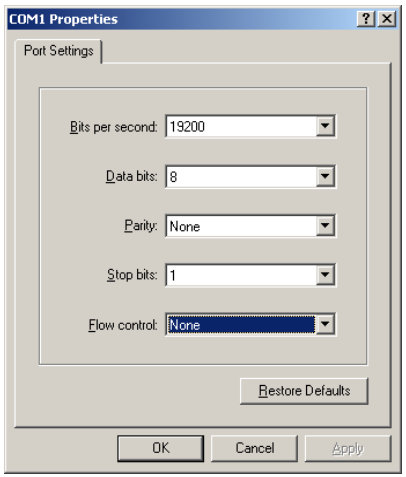

Figure 60: HyperTerminal COM1 Properties dialog box

e. In the **COM1 Properties** dialog box, enter the following settings, also listed in Figure 60, and click **OK**.

   Bits per Second: 19200
   Data bits: 8
   Parity: None
   Stop bits: 1
   Flow control: None

   The CLI input screen is displayed.

f. In the CLI session screen, press **Enter** several times to display the command prompt (`CLI>`).

   Commands can now be entered at this CLI prompt.

Help Commands

Help commands can be used to display a list of all possible commands or to display a detailed description of a specific command. Displayed information may include required command syntax, a brief definition, the number of characters for a command option value, or a list of allowed/disallowed characters.

- Displaying a List of All Basic Commands
- Displaying a List of All Available Commands
- Displaying Help for a Specific Command

The basic command is HELP, but can be modified with command options to more narrowly define the request.

Displaying a List of All Basic Commands

If the HELP command is entered without a limiting command verb or noun, the CLI will display the available basic commands.

Example Command

```
CLI> HELP
CLI> ?
```

Note: "HELP" and "?" are interchangeable.

Example Command Output

```
CLI> HELP

Possible command verbs:
                    help                add
delete              migrate             expand
extend              accept              rename
set                 locate              show

Possible command nouns:
unit                connection          acl
spare               units               unit_id
this_controller_id  other_controller_id globals
prompt              disk                bus
box                 all                 cancel
connections         this_controller     other_controller
version             disks               profile

Specify command word:
```

Displaying a List of All Available Commands

To see a list of all commands:

Example Command

```
CLI> HELP DISPLAY_ALL
```

Command Options

`DISPLAY_ALL`—a list of all currently supported commands will be displayed.

Example Command Output

```
CLI> HELP DISPLAY_ALL

Displaying list of all currently supported CLI commands:
?                           help
add unit                    add connection
add acl                     add spare
delete unit                 delete connection
delete acl                  delete spare
migrate unit                expand unit
extend unit                 accept unit
accept units                rename connection
set unit_id                 set this_controller_id
set other_controller_id     set globals
set acl                     set connection
set prompt                  set unit
locate disk                 locate unit
locate bus                  locate box
locate all                  locate cancel
show connections            show unit
show units                  show unit_id
show this_controller        show other_controller
show version                show disks
show globals                show acl
show profile

CLI>
```

Displaying Help for a Specific Command

When the basic HELP command is followed by a specific command verb or command verb and noun, the CLI will display helpful information about that command.

Example Command

```
CLI> HELP ADD ACL
```

Command Options

ADD ACL—the specific command about which to display help.

Example Command Output

```
CLI> HELP ADD

Possible command nouns:
unit                connection          acl
spare

Specify command noun:
```

Additional Command Example and Output

```
CLI> HELP ADD ACL

'add acl connection=name/wwpn=xxxxxxxx-xxxxxxxx unit=#'
adds access rights for one connection to one or all units.

CLI>
```

Display Commands

Several commands are available to view system information and setup, including:

- Viewing Physical Disk Information
- Viewing LUN Information
- Viewing LUN Names
- Viewing Controller Settings
- Viewing Global Settings
- Viewing Version Information

Each of these procedures uses the `SHOW` command, which displays the configuration of the MSA1000. These procedures are discussed in the following paragraphs.

Viewing Physical Disk Information

The following information can be displayed for the disks:

- Disk number
- Enclosure box and bay number
- Enclosure bus number and ID
- Disk size
- LUN in which the disk is used
- Disks assigned as spares

To display a list of physical disks contained in the MSA1000 and attached external storage enclosures:

Basic Command

```
SHOW DISKS
```

Example Command

```
CLI> SHOW DISKS
```

Example CLI Response

```
CLI> SHOW DISKS
 Disk List:   (box,bay)    (bus,ID)      Size        Units
   Disk101      (1,01)       (0,00)     18.2GB         0
   Disk102      (1,02)       (0,01)     18.2GB         0
   Disk103      (1,03)       (0,02)     18.2GB        none
   Disk104      (1,04)       (0,03)     18.2GB        none
   Disk105      (1,05)       (0,04)     18.2GB        none
   Disk106      (1,06)       (0,05)     18.2GB        none
   Disk107      (1,07)       (0,08)     18.2GB        none
   Disk108      (1,08)       (1,00)     18.2GB        none
   Disk109      (1,09)       (1,01)     18.2GB        none
   Disk110      (1,10)       (1,02)     18.2GB        none
   Disk111      (1,11)       (1,03)     18.2GB        none
   Disk112      (1,12)       (1,04)     18.2GB        none
   Disk113      (1,13)       (1,05)     18.2GB        none
   Disk114      (1,14)       (1,08)     18.2GB        none

CLI>
```

Viewing LUN Information

The following information can be displayed for the LUN:

- Unit identifier (user-defined name)
- LUN status, including OK, Failed, and Rebuilding
- List of the disks included in the LUN
- List of the disks assigned as spares to the LUN
- RAID level
- Unit size, in MB

Basic Command

```
SHOW UNIT
```

Example Command

To display all defined LUNs:

```
CLI> SHOW UNITS
```

To display a specific LUN:

```
CLI> SHOW UNIT 1
```

Example CLI Response

```
Unit 1:
In PDLA mode, Unit 1 IS LUN 2; In VSA mode, Unit 1 is LUN 1
Unit Identifier:
Device Indentifier:600805F3-00006B20-AE277D4B-B0D100F7
Cache Status:Enabled
Max Boot Partition:Disabled
Volume Status: VOLUME OK
Parity Init Status: 3% complete
5 Data Disk(s) used by lun 1:
  Disk107: Box 1, Bay 07,  (SCSI bus 0, SCSI id 8)
  Disk108: Box 1, Bay 08,  (SCSI bus 1, SCSI id 0)
  Disk207: Box 2, Bay 07,  (SCSI bus 2, SCSI id 8)
  Disk208: Box 2, Bay 08,  (SCSI bus 2, SCSI id 9)
  Disk307: BOX 3, Bay 07,  (SCSI bus 3, SCSI id 8)
  Disk308: Box 3, Bay 08,  (SCSI bus 3, SCSI id 9)
Spare Disk(s) used by lun 1:
  No spare drive is designated.
Logical Volume Raid Level: DISTRIBUTED PARITY FAULT TOLERANCE
(RAID 5)
                  stripe_size=16KB
Logical Volume Capacity: 173,658MB
```

Note: PDLA and VSA addressing method reference information is included in the display. PDLA represents Peripheral Device LUN Addressing, which is used by the majority of system profiles. VSA represents Volume Set Addressing, which is used by the HP profile.

Viewing LUN Names

To view the user-defined names that have been assigned to the LUNs:

Basic Command

```
SHOW UNIT_ID
```

Example Command

```
CLI> SHOW UNIT_ID 0
```

Command Options

`0`–the number of the LUN to display.

Example CLI Response

```
Unit 0: ABC
```

ABC represents the assigned user-defined name.

Viewing Controller Settings

The following information can be displayed for the controllers:

- Controller identifier
- Controller hardware and software versions
- SCSI compliance information
- Controller SCSI ID
- Redundancy information
- Host Port Information
- Controller cache configuration and battery information

To view information about the configuration of the controller:

Basic Command:

```
SHOW THIS_CONTROLLER
SHOW OTHER_CONTROLLER
```

Example Command

```
CLI> SHOW THIS_CONTROLLER
```

Example CLI Response

```
Controller:
  MSA1000© Compaq xxxxxxxxxxxx Software 2.30 Build 96
Hardware 7
  Controller Identifier: 123
  NODE_ID=yyyyyyyy-yyyyyyyy
  SCSI_VERSION=SCSI-3
  Supported Redundancy Mode: Active/Standby
  Current Redundancy Mode: Active/Standby
  Current Role:Active
  Device Port SCSI address 6
Host Port_1:
  REPORTED PORT_ID YYYYYYYY-YYYYYYYY
  PORT_1_TOPOLOGY=F_Port
Cache:
  128 megabytes read cache  128 megabytes write cache  Version 2
  Cache is GOOD, and Cache is enabled
  No unflushed data in cache
Battery:
  Module #1 is fully charged and turned off.
```

Viewing Global Settings

The following information can be displayed for the global settings of the controller:

- System name
- Rebuild and expand priority settings
- Read and write cache settings
- Power supply and EMU temperatures

To view information about the global settings of the controller:

Basic Command:

```
SHOW GLOBALS
```

Example Command

```
CLI> SHOW GLOBALS
```

Example CLI Response

```
CLI> SHOW GLOBALS

Global Parameters:
  System Name:  ABC
  Rebuild Priority:  high
  Expand Priority:   high

Total Cache:   512MB
  50% Read Cache:  256 MB
  50% Write Cache: 256 MB

Temperature:
  EMU:  29 Celsius,  84 Farenheit
  PS1:  40 Celsius  104 Farenheit
  PS2:  42 Celsius   107 Farenheit
```

Viewing Version Information

The following information can be displayed for the controllers:

- Firmware version
- Hardware revision
- Internal EMU revision

To view version information about the MSA1000:

Basic Command:

```
SHOW VERSION
```

Example Command

```
CLI> SHOW VERSION
```

Example CLI Response

```
CLI> SHOW VERSION
Firmware version:  2.28 build 80
Hardware revision:  7
Internal EMU Rev:   1.86
```

Array Controller Configuration Commands

The CLI provides the ability to configure the MSA1000 array controller, including:

- Setting Global Variables
- Setting the Controller ID
- Changing the CLI Prompt

Each of these procedures uses the `SET` command, which changes the settings of the MSA1000 Controller.

Setting Global Variables

A single command with a variety of parameters is used to set the expand priority, the read/write ratio, and the system name.

Basic Command

```
SET GLOBALS
```

Example Command

```
SET GLOBALS EXPAND_PRIORITY=HIGH REBUILD_PRIORITY=HIGH
SYSTEM_NAME="XXX" READ_CACHE=50 WRITE_CACHE=50
```

Command Options

`EXPAND_PRIORITY=HIGH`—the expand priority. Used when expanding an array to set the priority of array expansions in relation to input/output operations. See below for detailed information about these settings.

`REBUILD_PRIORITY=HIGH`—the rebuild priority. Used when rebuilding an array to set the priority of an array rebuild in relation to input/output operations. See below for detailed information about these settings.

Low—expansion or rebuild takes place only when the array controller is not busy handling normal I/O requests. This setting has minimal effect on normal I/O operations. However, there is an increased risk that data will be lost if another physical drive fails while the rebuild is in progress.

Medium—expansion or rebuild occurs for half of the time, and normal I/O requests are handled during the rest of the time.

High—rebuild or expansion occurs at the expense of normal I/O operations. Although system performance is affected, this setting provides better data protection because the array is vulnerable to additional drive failures for a shorter time.

SYSTEM_NAME="*XXX*"—the system name, where *XXX* represents any user defined phrase, up to 20 alphanumeric characters.

READ_CACHE=50—the read cache. This value must be between 0 and 100.

WRITE_CACHE=50—the write cache. This value must be between 0 and 100.

Note: Read_cache plus write_cache must equal 100.

Example CLI Response

```
Global Parameters:

System Name:XXX
Rebuild Priority:HIGH
Expand Priority:HIGH

Total Cache:128MB
50% Read Cache:64MB
50% Write Cache:64MB
```

Setting the Controller ID

To assign a unique name to each controller installed in the MSA1000:

Basic Command

```
SET THIS_CONTROLLER_ID
SET OTHER_CONTROLLER_ID
```

Example Command

```
CLI> SET THIS_CONTROLLER_ID 123
```

Command Options

123—the user-defined ID for the controller. The ID can be up to 230 alphanumeric characters or a decimal number in the range of 0 though 65535. In this example, the controller ID will be changed to 123.

Example CLI Response

```
CLI> SET THIS_CONTROLLER_ID 123
Controller identifier 123 created.
```

Note: To verify the controller ID assignment, use the SHOW THIS_CONTROLLER or SHOW OTHER_CONTROLLER commands.

Changing the CLI Prompt

To temporarily change the CLI prompt from the default of "CLI>":

Basic Command

```
SET PROMPT
```

Example Command

```
CLI> SET PROMPT ABC
```

Command Options

ABC—the user-defined name for the prompt. The name can be up to 24 alphanumeric characters. In this example, the prompt will be changed to ABC.

Example CLI Response

```
ABC>
```

LUN Management Commands

The following commands are used to physically locate the specific hard drives in a LUN as well as create, delete, and modify storage LUNs.

- Flashing LEDs/Locating Hard Drives
- Creating LUNs
- Assigning a Name or ID to a LUN
- Adding a Spare to a LUN
- Deleting LUNs
- Deleting Spares
- Recognizing a Failed Unit
- Modifying Arrays and LUNs

Flashing LEDs/Locating Hard Drives

A variety of commands is available for physically locating specific hard drives. Specifically, you can locate all drives attached to the MSA1000, all drives in a specific storage enclosure, all drives on a specific SCSI bus, all drives in a specific LUN, or a specific individual hard drive.

When these commands are executed, the LEDs of the requested drives will blink. These LEDs are visible from the front of the MSA1000 and its attached storage enclosures.

Note: If a time limit is not included with the Locate command, the LEDs will blink for 30 seconds.

The basic command verb is `LOCATE`, but a variety of command nouns and command options are available to customize the request.

Basic Command

```
LOCATE
```

Command Options

`TIME=xxx`—(optional) the length of time to flash the LEDs, where *xxx* represents the number of seconds.

`ALL`—all drives connected to the MSA1000 storage sub-system will be flashed.

`BOX`—the number of the storage enclosure whose disks you want to flash.

`1`=the MSA1000 drive shelf
`2`=the storage enclosure attached to SCSI port A
`3`=the storage enclosure attached to SCSI port B

`BUS`—the number of the bus whose disks you want to flash.

`UNIT`—the number of the LUN whose disks you want to flash.

`DISKxxx-DISKyyy`—a range of drives to flash. Disks are identified by box number and bay number.

`DISKzzz`—a specific drive to flash. Disks are identified by box number and bay number.

Locating all Drives Attached to the MSA1000

To flash the LEDs of all drives attached to the MSA1000:

Example Command

```
CLI> LOCATE ALL
```

Locating all Drives in a Specific Storage Enclosure

To flash the LEDs of all drives in the MSA1000 drive shelf or additional storage enclosures:

Example Command

```
CLI> LOCATE BOX 3
```

`BOX 3`—the drives in the storage enclosure attached to SCSI port B will be flashed.

Locating all Drives for a Specific SCSI Bus

To flash the LEDs of all drives connected to a specific SCSI bus:

Example Command

```
CLI> LOCATE BUS 1
```

`BUS 1`—all drives connected to bus number 1 will be flashed.

Locating all Drives in a Specific LUN

To flash the LEDs and locate all drives associated with a specific LUN:

Example Command

```
CLI> LOCATE UNIT 1
```

`UNIT 1`—the drives in LUN 1 will be flashed.

`UNIT ALL`—all drives incorporated into LUNs will be flashed.

Locating Specific Drives

To flash the LEDs to locate an individual drive or group of drives:

Example Command

```
CLI> LOCATE DISK DISK102-DISK106
```

`DISK DISK102-DISK106`—the LEDs on drives in bays 2 through 6 in box 1 will blink.

Canceling a LOCATE Request

To stop the flashing of the drive LEDs caused by a previous `Locate` command and return the drive LEDs to normal operations:

Example Command

```
CLI> LOCATE CANCEL
```

There are no available options for this command.

Creating LUNs

A LUN is a logical storage unit comprised of one or more hard drives.

When a LUN is initially created, the LUN Unit ID is automatically set to the number assigned to the LUN. If desired, this LUN ID name can be changed. See "Assigning a Name or ID to a LUN" for procedural information.

The basic command is `ADD UNIT`. A variety of options are available to modify the command.

Basic Command

```
ADD UNIT
```

Command Options

`DATA="DISKzzz"`—the number of an individual drive to incorporate into the LUN. Disks are identified by box number and bay number. For example, DISK110 identifies disk 10 in box number 1.

`DATA="DISKxxx-DISKyyy"`—a range of drives is to be incorporated into the LUN. Disks are identified by box number and bay number. For example, DISK101-DISK105 identifies disks 1 through 5 in box number 1.

Note: Quotation marks must be entered both before and after the data disk drives that are to be included in the LUN.

`RAID_LEVEL=`*`xxx`*—the RAID fault-tolerance level to use, where *`xxx`* represents:

`0` = RAID 0 (no fault tolerance)
`1` = RAID 1 (mirroring)
`5` = RAID 5 (distributed parity)
`ADG` = Advanced Data Guarding (ADG)

Note: If more than one pair of drives are included in a RAID 1 array, the data is striped across the first half of the drives in the array and then each drive is mirrored to a drive in the remaining half of the drives for fault tolerance. This method is referred to as RAID 1+0.

STRIPE_SIZE=*xxx*—(optional) the stripe size to assign (in KB), where *xxx* represents 8, 16, 32, 64, 128, or 256 Kilobytes.

RAID 0 uses stripe sizes 8, 16, 32, 64, 128, and 256 (Default: 128 KB)
RAID 1 uses stripe sizes 8, 16, 32, 64, 128, and 256 (Default: 128 KB)
RAID 5 uses stripe sizes 8, 16, 32, and 64 (Default: 16 KB)
RAID ADG uses stripe sizes 8, 16, 32, and 64 (Default: 16 KB)

SIZE=*xxxxyy*—(optional) how much of the available space on the indicated drives is to be used for the LUN, where *xxxx* represents the LUN size and *yy* indicates MB or GB. (When GB is entered, all return displays will be converted to MB.) If no size is specified, the maximum available space of the included disks will be assigned to the unit. See "Creating Multiple LUNs on a Group of Drives" for an example of using the Size command option.

SPARE=DISK*xxx*—(optional) disk(s) to assign as a spare to the unit, where *xxx* represents the disk number. More than one disk can be assigned as a spare to a LUN.

MAXBOOT=ENABLE/DISABLE—(optional) changes the size of the boot partition.

Enable = 8-GB boot partition (default)
Disable = 4-GB boot partition

CACHE=ENABLE/DISABLE—(optional) determines whether to use the array controller cache for the LUN.

Enable = (default) use the array controller cache
Disable = not use the array controller cache

Creating a Single LUN from a Group of Drives

To create one LUN from an individual drive or group of drives:

Example Command

```
CLI> ADD UNIT 0 DATA="DISK101-DISK107 DISK110" RAID_LEVEL=ADG
STRIPE_SIZE=64
```

Note: Maintain a record of the units as they are created. These Unit ID numbers are used in other CLI commands. In addition to recording the unit number, the drives included, RAID type and size, *record the order in which they are created*.

Command Options

0—LUN 0 is to be created.

DATA="DISK101-DISK107 DISK110"—drives 1 through 7 and disk 10 in box number 1 are to incorporated into the LUN.

RAID_LEVEL=ADG—RAID level ADG will be used when creating LUN 0.

STRIPE_SIZE=64—stripe size of 64 will be used when creating LUN 0.

Example CLI Response

```
First volume to be configured on these drives.
Logical Unit size = 69460 MB
RAID overhead = 0 MB
Total space occupied by new unit = 69460 MB
Free space left on this volume = 0 MB
Unit 0 is created successfully.
```

Creating Multiple LUNs on a Group of Drives

To carve out multiple LUNs from a drive or group of drives, use the `ADD UNIT` command and include the `SIZE=` command option. Repeat the command using a unique LUN ID for each LUN, along with the desired size parameter.

Example Command

```
CLI> ADD UNIT 1 DATA="DISK111-DISK114" RAID_LEVEL=5
STRIPE_SIZE=32 SIZE=1000MB
```

Note: Maintain a record of the units as they are created. These Unit ID numbers are used in other CLI commands. In addition to recording the unit number, the drives included, RAID type and size, and *record the order in which they are created*.

Command Options

`1`—LUN 1 is to be created.

`DATA="DISK111-DISK114"`—drives 11 through 14 in box number 1 to the LUN.

`RAID_LEVEL=5`—RAID 5 will be used when creating the LUN.

`STRIPE_SIZE=32`—stripe size of 32 will be used when creating the LUN.

`SIZE=1000MB`—1000 MB of the available space will be used for the LUN.

Example CLI Response

The following display is a continuation of the previous example, creating three LUNs on the same group of physical drives.

```
CLI> ADD UNIT 1 DATA="DISK111-DISK114" RAID_LEVEL=5
STRIPE_SIZE=32 SIZE=1000MB
First volume to be configured on these drives.
The logical unit size has been adjusted by 4MB for optimal
performance.
Logical Unit size = 996 MB
RAID overhead = 498 MB
Total space occupied by new unit = 1494 MB
Free space left on this volume = 24533 MB
Unit 1 is created successfully.

CLI> ADD UNIT 2 DATA="DISK111-DISK114" RAID_LEVEL=5
STRIPE_SIZE=32 SIZE=2000MB
Logical Unit size = 2000 MB
RAID overhead = 1000 MB
Total space occupied by new unit = 3000 MB
Free space left on this volume = 21533 MB
Unit 2 is created successfully.

CLI> ADD UNIT 3 DATA="DISK111-DISK114" RAID_LEVEL=5
STRIPE_SIZE=16 SIZE=4000MB
Logical Unit size = 4000 MB
RAID overhead = 2000 MB
Total space occupied by new unit = 6000 MB
Free space left on this volume = 15533 MB
Unit 3 is created successfully.
```

Creating a LUN with an Assigned Spare

To create a LUN and assign a spare at the same time, use the `ADD UNIT` command and include the `SPARE=` command option.

Note: More than one spare can be assigned to the same LUN and the same spare can be available to multiple LUNs.

Example Command

```
CLI> ADD UNIT 4 DATA="DISK211-DISK212" RAID_LEVEL=1
SPARE="DISK213"
```

Note: Maintain a record of the units as they are created. These Unit ID numbers are used in other CLI commands. In addition to recording the unit number, the drives included, RAID type and size, and record *the order in which they are created*.

Command Options

`4`—LUN 4 is to be created.

`DATA="DISK211-DISK212"`—drives 11 through 12 in box number 2 are to be incorporated into the LUN.

`RAID_LEVEL=1`—RAID 1 will be used when creating the LUN.

`SPARE="DISK213"`—the drive in bay 13 of box 2 will be assigned as a spare drive to the LUN.

Note: Quotation marks must be entered both before and after the disk drive that is to be designated as a spare for the LUN.

Example CLI Response

```
First volume to be configured on these drives.
Logical Unit size = 69460 MB
RAID overhead = 69460 MB
Total space occupied by new unit = 138920 MB
Free space left on this volume = 0 MB
Unit 4 is created successfully.
```

Assigning a Name or ID to a LUN

If desired (or required by your operating system), each LUN can be assigned a unique name or ID in addition to its number. These user-defined names make it easier to identify specific LUNs in other configuration procedures.

Note: OpenVMS systems require each LUN to have a unique ID. No two devices in the entire SAN for this OpenVMS system may share ID numbers. LUNs in different storage systems must have different IDs.

Basic Command

```
SET UNIT_ID
```

Example Command

```
CLI> SET UNIT_ID 0 ABC
```

Command Options

`0`–LUN 0 is being assigned a name.

`ABC`—is the name to assign to LUN 0. The name of this LUN will be ABC.

Example CLI Response

```
Identifier "ABC" created for unit 0
```

Adding a Spare to a LUN

To add a spare to an existing LUN:

Basic Command

```
ADD SPARE
```

Example Command

```
CLI> ADD SPARE UNIT=2 DISK109
```

Command Options

`UNIT=2`—the unit to assign the spare to. This is the same number that was given to the unit when it was created with the `ADD UNIT` command. In this example, LUN 2 will have use of the spare.

`DISK109`—indicates the drives to assign as the spare to the LUN. Disks are identified by box number and bay number. In this example, drive 9 in box number 1 is to be used.

Example CLI Response

```
First volume to be configured on these drives.
Logical Unit size = 69460 MB
RAID overhead = 69460 MB
Total space occupied by new unit = 138920 MB
Unit 2 is created successfully.
```

Deleting LUNs

To delete the last created LUN:

Note: If more than one LUN has been created, only the last LUN created can be deleted. It is important to maintain a record of the unit numbers and the order in which they are created.

Note: After a LUN is deleted, its unit number goes unused until manually assigned to a new LUN. Unit numbers are not automatically reassigned when a LUN is deleted.

Basic Command

```
DELETE UNIT
```

Example Command

```
CLI> DELETE UNIT 4
```

Command Options

4—the unit to delete. This is the same number that was given to the unit when it was created with the `ADD UNIT` command. In this example, LUN 4 is to be deleted.

Example CLI Response

```
Data will be lost after the unit is deleted.
Do you still want to DELETE unit 4 (Y/N)? Y
Please wait while unit 4 is being deleted…
Unit 4 is deleted successfully.
```

Deleting Spares

To remove a spare from use:

Basic Command

```
DELETE SPARE
```

Example Command

```
CLI> DELETE SPARE UNIT=2 DISK109
```

Command Options

`UNIT 2`—the unit that will no longer have access to the spare. This is the same number that was given to the unit when it was created with the `ADD UNIT` command. In this example, LUN 2 will no longer have access to the spare.

`DISK109`—the spare drive to remove from use. Disks are identified by box number and bay number. In this example, drive 1 in box number 9 will no longer be used as a spare for LUN 2.

Recognizing a Failed Unit

If all drives of a previously failed unit are in working order, use this command to change the state of the unit back to `VOLUME_OK`.

To accept media exchange on a unit marked as failed:

Basic Command

```
ACCEPT UNIT
ACCEPT UNITS
```

Example Command

```
CLI> ACCEPT UNIT 2
```

Command Options

`#`—the unit that you want to activate, where # represents the unit number.

If a unit number is not specified, all units will be reset.

Modifying Arrays and LUNs

Occasionally, after an array or a LUN has been created, its characteristics need to be changed. The following changes can be made:

- Adding Drives to an Array
- Adding Space to a LUN
- Changing the RAID Characteristics of a LUN
- Changing Attributes of a LUN

Each task is discussed in the following paragraphs.

Adding Drives to an Array

To add additional physical drives to an array:

Note: Because this command affects the entire array, all LUNs made from the array are also affected.

Basic Command

```
EXPAND UNIT
```

Example Command

```
CLI> EXPAND UNIT 4 DISK204-DISK207
```

Command Options

`4`—one of the LUN within the target array.

Note: Any LUN in the array can be entered to identify the array; the space is added to the array and not the LUN.

`DISK204-DISK207`—the physical disks to add to the array. Disks are identified by box number and bay number. In this example, `DISK204-DISK207` identifies disks 4 through 7 in box number 2.

Example CLI Response

```
The actual new array capacity will be 3000MB.
The array with Unit 4 is being expanded.
Use "show unit 4" to monitor progress.
```

Adding Space to a LUN

To add unused, available space in an array to a specific LUN:

Basic Command

```
EXTEND UNIT
```

Example Command

```
CLI> EXTEND UNIT 2 ADD_SIZE=1000MB
```

Command Options

`2`—the LUN to which the space will be added. In this example, LUN 2 is being extended.

`ADD_SIZE=1000MB`—how much of the available space in the array to add to the LUN. In this example, 1000 MB of space will be added. The size limit must be specified as GB, MB, or KB. If no size is specified, the maximum available space of the included disks will be assigned to the array.

`NEW_SIZE=`*xxxxyy*—can be used instead of `ADD_SIZE` to enter the total new size of the LUN, where *xxxx* represents the size and *yy* indicates GB, MB, or KB.

Example CLI Response

```
The actual new volume size will be 1992MB.
Unit 2 is being extended.
Use "show unit 2" to monitor progress.
```

Changing the RAID Characteristics of a LUN

While obeying the constraints of included number of drives and applicable stripe sizes, you can migrate an array from one RAID level to another. If you attempt to move to an unsupported RAID configuration for an array, an error message is displayed.

To change the RAID level or the stripe size of a LUN:

Note: Before changing the RAID level or stripe size of a LUN, verify there is available, unused space on the array. Migrating from one RAID level to another may require additional space for parity and organizational purposes.

Basic Command

```
MIGRATE UNIT
```

Example Command

```
CLI> MIGRATE UNIT 0 RAID_LEVEL=5 STRIPE_SIZE=32
```

Command Options

`UNIT 0`—the number of the LUN to modify. In this example, LUN 0 will be migrated.

`RAID_LEVEL=`*xxx*—RAID 5 will be assigned to the LUN.
RAID levels include:

`0` = RAID 0 (no fault tolerance)
`1` = RAID 1 (mirroring)
`5` = RAID 5 (distributed parity)
`ADG` = Advanced Data Guarding (ADG)

`STRIPE_SIZE=`*xxx*—(optional) the stripe size of 32 will be used. Other stripe size options include:

`RAID 0` uses stripe sizes 8, 16, 32, 64, 128, and 256 (Default: 128 KB)
`RAID 1` uses stripe sizes 8, 16, 32, 64, 128, and 256 (Default: 128 KB)
`RAID 5` uses stripe sizes 8, 16, 32, and 64 (Default: 16KB)
`RAID ADG` uses stripe sizes 8, 16, 32, and 64 (Default: 16 KB)

Example CLI Response

```
The RAID level of Unit 0 will now be 5.
Unit 0 is being migrated.
Use "show unit 0" to monitor progress.
```

Changing Attributes of a LUN

To enable or disable the array accelerator cache for a specific LUN:

Basic Command

```
SET UNIT
```

Example Command

```
CLI> SET UNIT 0 CACHE=DISABLE
```

Command Options

`UNIT 0`—the number of the LUN to modify. In this example, LUN 0 will be modified.

`CACHE=ENABLE or DISABLE`—either enables or disables the use of the array accelerator cache for the specified LUN.

Example CLI Response

```
CLI> SET UNIT 0 CACHE=ENABLE
Cache for unit 0 has been enabled.
```

```
CLI> SET UNIT 1 CACHE=DISABLE
Cache for unit 1 has been disabled.
```

Server Connection Commands

For a server to access the storage of the MSA1000, you must first establish a connection that identifies the specific HBA to the MSA1000.

After connections are entered, you may need to restrict access to LUNs to specific servers. Some programs, including the Array Configuration Utility (ACU), use the phrase Selective Storage Presentation (SSP) when discussing storage limitations. The CLI uses an Access Control List (ACL) to enter the list of LUNs a server can access. See the following section "Selective Storage Presentation/Access Control List Commands" for more information about SSP and ACLs.

Each of the following commands is used to enter and manage the connections:

- Viewing Existing Connections and the HBA World Wide Name
- Naming a Connection to a server
- Managing Connection Profiles

Viewing Existing Connections and the HBA World Wide Name

The following command displays the World Wide Names of each Host Bus Adapter (HBA) attached to the MSA1000. If connections between the HBAs and the MSA1000 have been given user-defined names, these names are also displayed.

Basic Command

```
SHOW CONNECTIONS
```

Example Command

```
CLI> SHOW CONNECTIONS
```

Command Options

`CONNECTION NAME`—the name of a specific connection to display.

`WWPN`—the WWPN of a specific connection to display.

Example CLI Response

```
Connection Name: abc
  Host WWNN = 11111111-11111111
  Host WWPN = 12345678-12345678
  Profile Name = Default
  Unit Offset = 0
  Controller 1 Port 1 Status = Online
  Controller 2 Port 1 Status = Online

Connection Name: <unknown>
  Host WWNN = 22222222-22222222
  Host WWPN = 33333333-33333333
  Profile Name = Default
  Unit Offset = 0
  Controller 1 Port 1 Status = Online
  Controller 2 Port 1 Status = Online
```

Naming a Connection to a server

This command is used to create an association between the Fibre Channel Initiator (HBA in the server) and the MSA1000, while at the same time naming and identifying the host mode (operating system) of the server.

The name given to a connection is user defined, and according to a naming convention that makes identification and setup of the ACL for this connection easier.

To name a connection between the MSA1000 and a server:

Basic Command

```
ADD CONNECTION
```

Example Command

```
CLI> ADD CONNECTION ABC WWPN=12345678-12345678 PROFILE=WINDOWS
```

Command Options

`ABC`—the name to give the connection.

`WWPN=12345678-12345678`—the WWPN of the active HBA inside the server attached to the MSA1000. The show connections command can be used to obtain the WWPN.

`WWNN`—the WWNN of the active HBA inside the server attached to the MSA1000. The show connections command can be used to obtain the WWNN.

`PROFILE=WINDOWS`—the platform of the host. If a profile is not specified, the default profile is used.

Profile options include:

Windows (default)
OVMS
Tru64
Linux
Solaris
NetWare
HP

`OFFSET=X`—(default: 0) the unit offset for assigning logical volumes.

Example CLI Response

```
Connection has been added successfully.
Profile Windows is set for the new connection.
```

Managing Connection Profiles

To change the operating-system profile associated with an existing connection:

Note: The connection profile is sometimes referred to as the Host Mode.

Basic Command

```
SET CONNECTION X PROFILE=Y
```

Example Command

```
CLI> SET CONNECTION ABC PROFILE=WINDOWS
```

Command Options

`ABC`—the name of the connection to modify, in this example ABC.

`WWPN=12345678-12345678`—the WWPN of the connection to modify, using the WWPN of the HBA. The "show connections" can be used to obtain the WWPN.

`WWNN=12345678-12345678`—the WWNN of the connection to modify, using the WWNN of the HBA. The "show connections" can be used to obtain the WWNN.

`PROFILE=WINDOWS`—the platform of the host, which in this example is Windows. See "Naming a Connection to a Server" for a list of profile types.

Example CLI Response

```
The Profile of Connection ABC is set to Windows successfully.
```

Selective Storage Presentation/Access Control List Commands

When multiple servers access the storage of the MSA1000, it becomes necessary to restrict access to LUNS to specific servers. Some programs, including the Array Configuration Utility (ACU), use the phrase “selective storage presentation” (SSP) when discussing storage limitations. The CLI uses an Access Control List (ACL) to enter the list of LUNs a server can access.

Each of the following commands is used to set up and manage the ACL:

- Viewing the Connection Names
- Viewing the ACL
- Adding to the ACL
- Deleting Information from the ACL
- Changing the Name of a Connection
- Changing the HBA of a Connection
- Deleting a Connection Name
- Disabling the ACL

Each of these tasks is defined in the following sections.

Note: Many commands are discussed in this section. To set up the ACL, the following commands should be used in sequence:

- Viewing the Connection Names
- Naming a Connection to a server
- Entering ACL information

Viewing the Connection Names

Use the SHOW CONNECTIONS command to display the connection name associated with each Host Bus Adapter (HBA) attached to the MSA1000.

The connection names are used when entering ACL information.

See "Viewing Existing Connections and the HBA World Wide Name" for details about this command.

Viewing the ACL

To display the current ACL:

Basic Command

```
SHOW ACL
```

Command Options

There are no available options for this command.

Example CLI Response

```
ACL is enabled:
Connection      WWPN                  Units
ABC             11111111-22222222     0,1,2
XYZ             33333333-44444444     2,3,4
Inaccessible Units: 5,6
```

Adding to the ACL

To indicate the allowable LUNs for each server:

Basic Command

```
ADD ACL
```

Command Options

CONNECTION=*xxx*—the name of the connection to grant access, where *xxx* represents the connection name.

UNIT=*xxx*—which LUN to assign to the indicated server.

Options include:

An individual LUN can be assigned by entering one LUN ID (UNIT=0)
A group of LUN can be assigned by entering a range of LUN (UNIT=1-3)
Access to all units can be granted by entering UNIT=ALL.

Two methods of entering LUN assignments are available:

- Adding to the ACL using the connection name
- Adding to the ACL using the WWPN

Note: There is no command to activate the ACL. Immediately after the first entry is added to the ACL, access to the storage is limited to the servers and LUNs listed in the ACL.

Adding to the ACL using the Connection Name

Example Command

```
CLI> ADD ACL CONNECTION=ABC UNIT=ALL
```

Command Options

`CONNECTION=ABC`—entries will be added to the ACL of the connection named ABC.

`UNIT=ALL`—all LUN will be accessible to the connection named ABC.

Example CLI Response

```
Allowing 12345678-12345678 access to unit 2.
```

Adding to the ACL using the WWPN

Example Command

```
CLI> ADD ACL WWPN=12345678-12345678 UNIT=2
```

Command Options

`WWPN=12345678-12345678`—entries will be added to the ACL of the connection with a WWPN of 12345678-12345678.

`UNIT=2`—LUN 2 will be accessible to the named connection.

Example CLI Response

```
Allowing 12345678-12345678 access to unit 2.
```

Deleting Information from the ACL

After the ACL has been set up, access permissions for certain LUNs may need to be removed.

Basic Command

```
DELETE ACL
```

Command Options

`CONNECTION=`*`xxx`*—the name of the connection for which to remove permissions.

`UNIT=`*`xxx`*—which LUNs to remove from the ACL of the indicated server.

Two methods of removing access to previously assigned LUNs are available:

- Deleting information from the ACL using the connection name
- Deleting information from the ACL using the WWPN

Note: If all entries in the ACL are being deleted, immediately after the last entry is deleted, all connected servers have unlimited access to the storage.

Deleting Information from the ACL using the connection name

Example Command

```
CLI> DELETE ACL CONNECTION=ABC UNIT 0
```

Command Options Used

`CONNECTION=ABC`—entries will be deleted from the ACL of the connection named ABC.

`UNIT=0`—LUN 0 is to be deleted from the ACL of the named connection.

Example CLI Response

```
Disallowing 12345678-12345678 access to unit 0.
```

Deleting Information from the ACL using the WWPN

Example Command

```
CLI> DELETE ACL WWPN=12345678-12345678 UNIT=1
```

Command Options

`WWPN=12345678-12345678`—entries will be deleted from the ACL of the connection with a WWPN of 12345678-12345678.

`UNIT=1`—LUN 1 will be removed from the ACL of the named connection.

Example CLI Response

```
Disallowing 12345678-12345678 access to unit 1.
```

Changing the Name of a Connection

To change the name associated with a connection:

Basic Command

```
RENAME CONNECTION
```

Example Command

```
CLI> RENAME CONNECTION ABC XYZ
```

Command Options

`ABC`—the current name of the connection that is to be changed.

`XYZ`—the new name to assign to the connection, up to 16 alphanumeric characters.

Example CLI Response

```
Connection(s) has been renamed successfully.
```

Changing the HBA of a Connection

When a new HBA needs to be associated with an existing connection name:

Basic Command

```
SET CONNECTION
```

Example Command

```
CLI> SET CONNECTION ABC WWPN=12345678-Y999999Y
```

Command Options

`ABC`—the connection to make the changes to.

`WWPN=12345678-12345678`—the World Wide Port Name (WWPN) of the new HBA to associate with the connection.

`WWNN=12345678-12345678`—the World Wide Node Name (WWNN) of the new HBA to associate with the connection.

Example CLI Response

```
Connection(s) has been set successfully.
The WWPN of connection ABC is set to 12345678-Y999999Y
successfully.
```

Deleting a Connection Name

To remove the name associated with a connection to a server:

Basic Command

```
DELETE CONNECTION
```

Example Command

```
CLI> DELETE CONNECTION ABC
```

Command Options

`ABC`—the nickname that was assigned to the HBA within the server.

Example CLI Response

```
Connection(s) has been deleted successfully.
```

Disabling the ACL

If the ACL has been set up, but the decision has been to not use it and not limit access to the storage, the entries in the ACL need to be deleted. Use the following command to delete all entries in the ACL.

Basic Command

```
SET ACL DISABLE
```

Example CLI Response

```
Disabling acl
```

Note: When this command is used, all existing ACL entries are automatically deleted. All storage is immediately available to all connected servers.

Regulatory Compliance Notices

A

Regulatory Compliance Identification Numbers

For the purpose of regulatory compliance certifications and identification, your HP StorageWorks Modular SAN Array 1000 is assigned an HP Series number. The HP Series number for this product is: Series E01501. The Storage System Series number can be found on the product label, along with the required approval markings and information. The product label is located on the chassis and inside the fan module opening. When requesting certification information for this product, always refer to this Series number. This Series number should not be confused with the marketing name or model number for your Storage System.

Federal Communications Commission Notice

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at personal expense.

Modifications

The FCC requires the user to be notified that any changes or modifications made to this device that are not expressly approved by Hewlett-Packard Company may void the user's authority to operate the equipment.

Cables

Connections to this device must be made with shielded cables with metallic RFI/EMI connector hoods in order to comply with FCC Rules and Regulations.

Canadian Notice (Avis Canadien)

This Class A digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la classe A respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

European Union Notice

Products with the CE Marking comply with both the EMC Directive (89/336/EEC) and the Low Voltage Directive (73/23/EEC) issued by the Commission of the European Community.

Compliance with these directives implies conformity to the following European Norms (in brackets are the equivalent international standards):

- EN55022 (CISPR 22) - Electromagnetic Interference
- EN50082-1 (IEC801-2, IEC801-3, IEC801-4) - Electromagnetic Immunity
- EN60950 (IEC950) - Product Safety

Japanese Notice

ご使用になっている装置にVCCIマークが付いていましたら、次の説明文をお読み下さい。

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラスB情報技術装置です。この装置は、家庭環境で使用することを目的としていますが、この装置がラジオやテレビジョン受信機に近接して使用されると、受信障害を引き起こすことがあります。
取扱説明書に従って正しい取り扱いをして下さい。

VCCIマークが付いていない場合には、次の点にご注意下さい。

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラスA情報技術装置です　この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

BSMI Notice

警告使用者：

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

Laser Compliance

The SFP Module contains a laser diode of either gallium aluminum arsenide (GaALAs) emitting in the wavelength range of 770-860 nm, or indium gallium arsenide phosphide (InGaAsP) emitting in the wavelength range of 1270-1355 nm. All HP systems equipped with a laser device comply with safety standards, including International Electrotechnical Commission (IEC) 825. With specific regard to the laser, the equipment complies with laser product performance standards set by government agencies as a Class 1 laser product. The product does not emit hazardous laser radiation.

WARNING: Use of controls or adjustments or performance of procedures other than those specified herein or in the laser product's installation guide may result in hazardous radiation exposure. To reduce the risk of exposure to hazardous radiation:

- Do not try to open the unit enclosure. There are no user-serviceable components inside.
- Do not operate controls, make adjustments, or perform procedures to the laser device other than those specified herein.
- Allow only HP Authorized Service technicians to repair the unit.

The Center for Devices and Radiological Health (CDRH) of the U.S. Food and Drug Administration implemented regulations for laser products on August 2, 1976. These regulations apply to laser products manufactured from August 1, 1976. Compliance is mandatory for products marketed in the United States. This device is classified as a Class 1 laser product as defined by IEC 825.

This indicates that the product is classified as a CLASS 1 LASER PRODUCT.

Battery Replacement Notice

Your MSA1000 is provided with a Lithium Manganese Dioxide, or a Vanadium Pentoxide, batteries. There is a danger of explosion and risk of personal injury if the array is incorrectly replaced or mistreated. Replace only with the HP spare designated for this product. For more information about battery replacement or proper disposal, contact your HP Authorized Reseller or your Authorized Service Provider.

WARNING: Your Accelerator Array contains Lithium manganese Dioxide, or Vanadium Pentoxide batteries. There is risk of fire and burns if the battery pack is not handled properly. To reduce the risk of personal injury:

- Do not attempt to recharge the battery.
- Do not expose to temperatures higher than 60 C.
- Do not disassemble, crush, puncture, short external contacts, or dispose of in fire or water.
- Replace only with HP spare designated for this product.

Caution: Batteries, battery packs, and accumulators should not be disposed of together with the general household waste. In order to forward them to recycling or proper disposal, please use the public collection system or return them to HP, your authorized HP Partners, or their agents.

Electrostatic Discharge

B

To prevent damage to the system, be aware of the precautions you need to follow when setting up the system or handling parts. A discharge of static electricity from a finger or other conductor may damage system boards or other static-sensitive devices. This type of damage may reduce the life expectancy of the device.

To prevent electrostatic damage, observe the following precautions:

- Avoid hand contact by transporting and storing products in static-safe containers.
- Keep electrostatic-sensitive parts in their containers until they arrive at static-free workstations.
- Place parts on a grounded surface before removing them from their containers.
- Avoid touching pins, leads, or circuitry.
- Always make sure you are properly grounded when touching a static-sensitive component or assembly.

Grounding Methods

There are several methods for grounding. Use one or more of the following methods when handling or installing electrostatic-sensitive parts:

- Use a wrist strap connected by a ground cord to a grounded workstation or computer chassis. Wrist straps are flexible straps with a minimum of 1 megohm ± 10 percent resistance in the ground cords. To provide proper ground, wear the strap snug against the skin.
- Use heel straps, toe straps, or boot straps at standing workstations. Wear the straps on both feet when standing on conductive floors or dissipating floor mats.
- Use conductive field service tools.
- Use a portable field service kit with a folding static-dissipating work mat.

If you do not have any of the suggested equipment for proper grounding, have an Authorized HP Reseller install the part.

Note: For more information on static electricity, or assistance with product installation, contact your HP Authorized Reseller.

Specifications

C

This appendix provides operating and physical specifications for the MSA1000.

Table 17: MSA1000 Specifications

Parameter	English	Metric
Dimensions Height Depth Width	 6.9 in 20.5 in 19.0 in	 17.5 cm 52.1 cm 48.3 cm
Weight No drives installed, single power supply	50 lb	22.7 kg
Input Power Requirements Rated Input Voltage Rated Input Frequency Rated Input Current Input Power (max)	 100 to 240 VAC 50 - 60 Hz 7.35 A Max 641 W*	 100 to 240 VAC 50 - 60 Hz 7.35 A Max 641 W*
Heat Dissipation (max)	2187 Btu/hr*	2187 Btu/hr*
Temperature Range Operating Shipping	 50° to 95° F -22° to 122° F	 10° to 35° C (derated 1°C per 1000 feet of elevation to 10000 ft.) -30° to 50° C
Relative Humidity (non-condensing) Operating Non-operating	 10% to 90% up to 95%	 10% to 90% up to 95%

Table 17: MSA1000 Specifications

Parameter	English	Metric
Maximum Wet Bulb Temperature		
Long term storage	84.2°F	29°C
Short term storage	86°F	30°C
*Input Power and Heat Dissipation specifications are maximum values and apply to worst-case conditions at full rated power supply load. The power/heat dissipation for your installation will vary depending on the equipment configuration.		

Hard Drive Arrays

This appendix explains drive array concepts and data protection methods, including fault-tolerance options. This information will help you decide how to best configure your controller.

The following sections are included:

- What is a Drive Array?
- Fault-Tolerance Methods

What is a Drive Array?

Business users demand high storage capacities, high data transfer rates, and great protection against data loss when drives fail.

Connecting extra physical drives to a system increases the total storage capacity (refer to Figure 61), but has no effect on the efficiency of read/write (R/W) operations. Data can still be transferred to only one physical drive at a time.

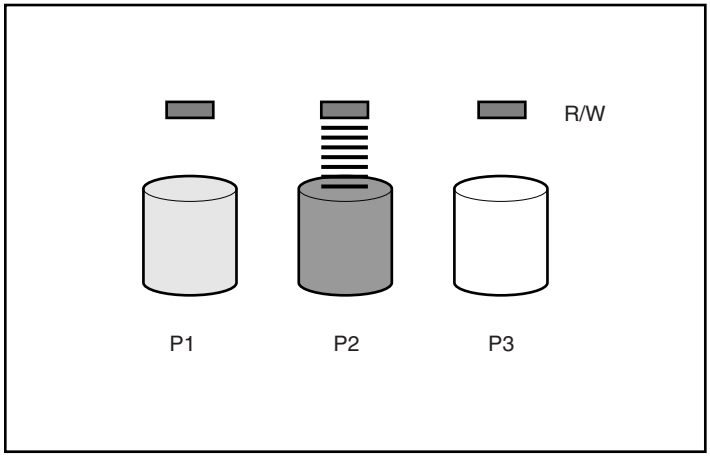

Figure 61: Physical drives added to system

With an array controller installed in the system, the capacity of several physical drives can be combined into one or more virtual units called **logical drives** (also called logical volumes).

In a logical drive, the read/write heads of all the constituent physical drives are active simultaneously, reducing the total time required for data transfer.

Figure 62 is an illustration of a logical drive spanning three physical hard drives.

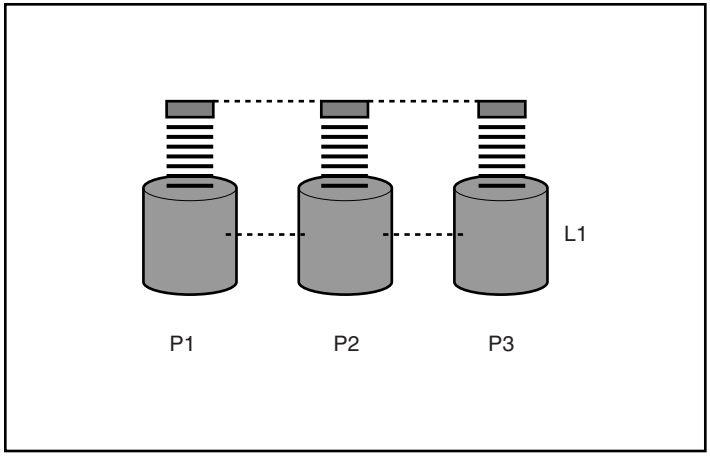

Figure 62: Physical drives configured into a logical drive (L1)

Because the read/write heads are active simultaneously, the same amount of data is written to each drive during any given time interval. Each unit of data is called a **block**, and over all the physical drives in a logical drive the blocks form a set of data **stripes** (refer to Figure 63).

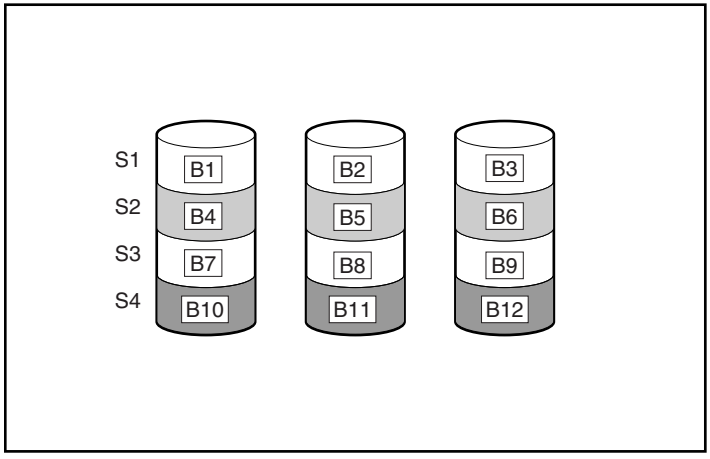

Figure 63: Data striping (S1-S4) of data blocks B1-B12

For data in the logical drive to be readable, the data block sequence must be the same in every stripe. This sequencing process is performed by the array controller, which sends the data blocks to the drive write heads in the correct order.

A natural consequence of the striping process is that each physical drive in a given logical drive will contain the same amount of data. If one physical drive has a larger capacity than other physical drives in the same logical drive, the extra capacity is wasted because it cannot be used by the logical drive.

The group of physical drives containing the logical drive is called a **drive array** (or just **array**). An array can contain several logical drives, each of a different size (refer to Figure 64.)

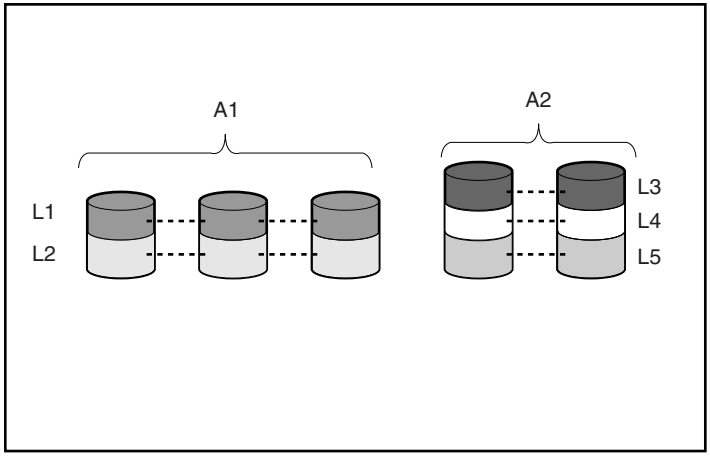

Figure 64: Two arrays (A1, A2) containing five logical drives

Each logical drive in an array is distributed over all of the physical drives within the array. A logical drive can also extend over more than one port on the same controller, but it cannot extend over more than one controller.

Drive failure, although rare, is potentially catastrophic. In Figure 64, for example, failure of **any** physical drive causes **all** logical drives in the same array to fail, and all data on the drives is lost.

To protect against data loss due to physical drive failure, logical drives are configured with **fault tolerance.** There are several fault-tolerance methods; those described in the following section are:

- RAID 0—Data Striping only (no fault tolerance)
- RAID 1—Drive Mirroring
- RAID 1+0—Data Striping and Drive Mirroring
- RAID 5—Distributed Data Guarding
- RAID ADG—Advanced Data Guarding

For any configuration except RAID 0, further protection against data loss can be achieved by assigning a drive as an **online spare** (or **hot spare**). This drive contains no data and is connected to the same controller as the array. When any other physical drive in the array fails, the controller automatically rebuilds information that was originally on the failed drive to the online spare. The system is quickly restored to full RAID-level data protection. (However, in the unlikely event that another drive in the array fails while data is being rewritten to the spare, the logical drive will still fail.)

When you configure an online spare, it is automatically assigned to all logical drives in the same array. Additionally, you do not need to assign a separate online spare to each array; you can configure one hard drive to be the online spare for several arrays, as long as the arrays are all on the same controller.

Fault-Tolerance Methods

RAID 0—No Fault Tolerance

This configuration (refer to Figure 63) provides no protection against data loss when a drive fails. However, it is useful for rapid storage of large amounts of non-critical data (for printing or image editing, for example) or when cost is the most important consideration.

Advantages

- Highest performance method for writes
- Lowest cost per unit of data stored
- All drive capacity is used to store data (none needed for fault tolerance)

Disadvantages

- All data on the logical drive is lost if a physical drive fails
- Cannot use an online spare
- Can only preserve data by backing it up to external drives

RAID 1 Drive Mirroring/RAID 1+0—Data Striping and Drive Mirroring

RAID 1 configurations contain only two drives. In this configuration, data is duplicated to a second drive.

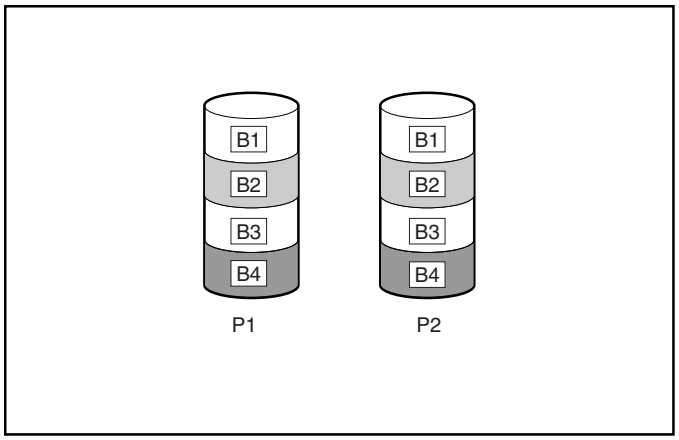

Figure 65: Drive mirroring of P1 to P2

When a RAID 1 array has more than two physical drives, this fault-tolerance method is referred to as RAID 1+0. In this configuration, the data blocks are striped across half of the drives in the array and then the drives are mirrored in pairs. See Figure 66 for an illustration of this configuration method.

In each mirrored pair, the physical drive that is not busy answering other requests answers any read request sent to the array. (This behavior is called **load balancing**.) If a physical drive fails, the remaining drive in the mirrored pair can still provide all the necessary data. Several drives in the array can fail without incurring data loss, as long as no two failed drives belong to the same mirrored pair.

This fault-tolerance method is useful when high performance and data protection are more important than the cost of physical drives.

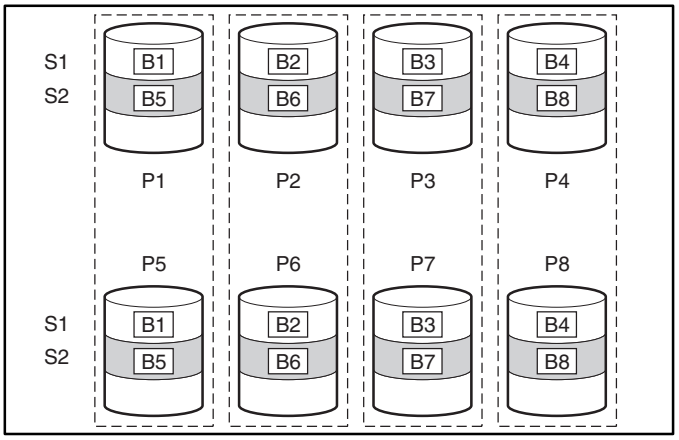

Figure 66: Data Striping and Drive Mirroring

Advantages

- Highest read and write performance of any fault-tolerant configuration
- No loss of data as long as no failed drive is mirrored to another failed drive (up to half of the physical drives in the array can fail)

Disadvantages

- Expensive (many drives needed for fault tolerance)
- Only half of total drive capacity usable for data storage

RAID 5—Distributed Data Guarding

By this method, a block of **parity data** is calculated for each stripe from the data that is in all other blocks within that stripe. The blocks of parity data are distributed over every physical drive within the logical drive (refer to Figure 67). When a physical drive fails, data that was on the failed drive can be calculated from the user data on the remaining drives and the parity data. This recovered data is usually written to an online spare in a process called a **rebuild.**

This configuration is useful when cost, performance, and data availability are equally important.

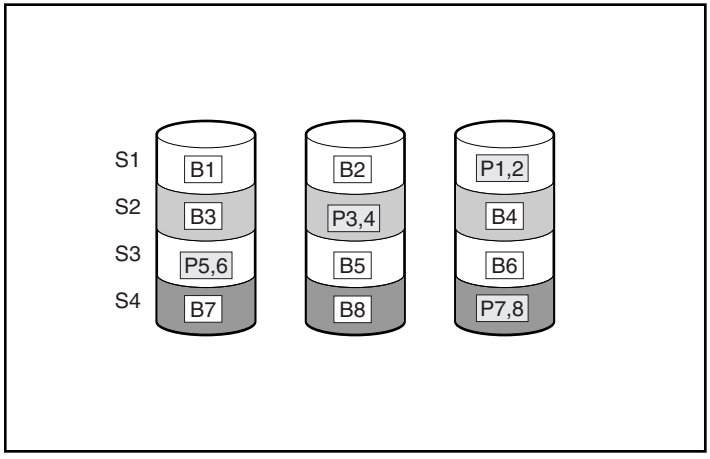

Figure 67: Distributed data guarding, showing parity information (Px,y)

Advantages

- High read performance
- No loss of data if one physical drive fails
- More drive capacity usable than with RAID 1+0—parity information requires only the storage space equivalent to one physical drive

Disadvantages

- Relatively low write performance
- Loss of data if a second drive fails before data from the first failed drive is rebuilt

RAID ADG—Advanced Data Guarding

RAID ADG is similar to RAID 5 in that parity information is generated (and stored) to protect against data loss caused by drive failure. With RAID ADG, however, two different sets of parity data are used, allowing data to still be preserved if two drives fail. As can be seen in Figure 68, each set of parity data uses up a capacity equivalent to that of one of the constituent drives.

This method is most useful when data loss is unacceptable, but cost must also be minimized. The probability that data loss will occur when arrays are configured with RAID ADG is less than when they are configured with RAID 5.

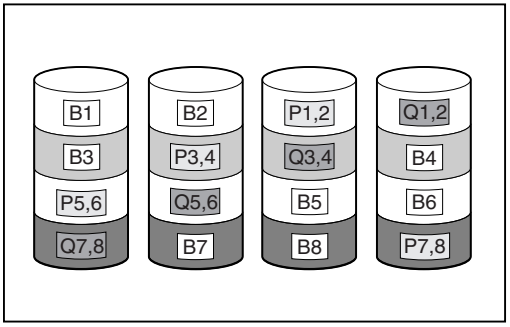

Figure 68: Advanced data guarding (RAID ADG)

Advantages

- High read performance
- High data availability—any two drives can fail without loss of critical data
- More drive capacity usable than with RAID 1+0—parity information requires only the storage space equivalent to two physical drives

Disadvantage

The only significant disadvantage of RAID ADG is a relatively low write performance (lower than RAID 5), because of the need for two sets of parity data.

Table 18 summarizes the important features of the different kinds of RAID methods described here. The decision chart in Table 19 may help you to determine which option is best for your situation.

Table 18: Summary of RAID Methods

	RAID 0	RAID1+0	RAID 5	RAID ADG*
Alternative name	Striping (no fault tolerance)	Striping and Mirroring	Distributed Data Guarding	Advanced Data Guarding
Usable drive space**	100%	50%	67% to 93%	50% to 96%
Usable drive space formula	n	n/2	(n-1)/n	(n-2)/n
Minimum number of physical drives	1	2	3	4
Tolerates physical drive failure?	No	Yes	Yes	Yes
Tolerates simultaneous failure of more than one physical drive?	No	Only if no two failed drives are in a mirrored pair	No	Yes
Read performance	High	High	High	High
Write performance	High	Medium	Low	High
Relative cost	Low	High	Medium	Medium

* Not all controllers support RAID ADG.

** Values for usable drive space are calculated with these assumptions: (1) all physical drives in the array have the same capacity; (2) online spares are not used; (3) no more than 14 physical drives are used per array for RAID 5; (4) no more than 56 drives are used with RAID ADG.

Table 19: Choosing a RAID Method

Most Important	Also Important	Suggested RAID Level
Fault tolerance	Cost effectiveness	RAID ADG
	I/O performance	RAID 1+0
Cost effectiveness	Fault tolerance	RAID ADG*
	I/O performance	RAID 5 (RAID 0 if fault tolerance is not required)
I/O performance	Cost effectiveness	RAID 5 (RAID 0 if fault tolerance is not required)
	Fault tolerance	RAID 1+0
*Not all controllers support RAID ADG		

Other Fault-Tolerance Options

The operating system that you are using may also support software-based RAID or controller duplexing.

- **Software-based RAID** resembles hardware-based RAID, except that the operating system works with logical drives as if they were physical drives. To protect against data loss caused by physical drive failure, each logical drive must be in a different array from the others.
- **Controller Duplexing** uses two identical controllers with independent, identical sets of drives containing identical data. In the unlikely event of a controller failure, the remaining controller and drives will service all requests.

However, the hardware-based RAID methods described in this appendix provide a much more robust and controlled fault-tolerant environment. Additionally, controller duplexing and software-based RAID do not support online spares, auto-reliability monitoring, interim data recovery, or automatic data recovery.

If you decide to use one of these alternative fault-tolerance options, configure your arrays with RAID 0 for maximum storage capacity and refer to your operating system documentation for further implementation details.

Recovering from Hard Drive Failure

E

Caution: Back up all data before removing drives or changing configurations. Failure to do so could result in permanent loss of data. Before moving drives and arrays, run the Array Configuration Utility.

The purpose of fault-tolerant configurations on the MSA1000 Controller is to protect against data loss due to drive failure. Although the MSA1000 Controller firmware is designed to protect against normal drive failure, it is imperative that you perform the correct actions to recover from a drive failure without inadvertently inducing any additional drive failures. Multiple drive failures in the same array generally result in data loss (exceptions include failures following activation of a spare drive, and failure of drives in a mirroring configuration that are not mirrored to one another or Advanced Data Guarding configuration).

Drives may be either Ultra2 or Ultra3 Pluggable Universal Hard Disk Drives. Drives that are to be grouped in the same array should be of the same capacity. The excess capacity of a larger drive cannot be used by the array and will be wasted.

A list of supported hard drives is available at http://www.hp.com.

Hard Drive Failure

When a hard drive fails, all logical drives that are in the same array will be affected. Each logical drive in an array may be using a different fault tolerance method, so each logical drive can be affected differently.

- RAID 0 configurations cannot tolerate drive failure. If any physical drive in the array fails, all non-fault-tolerant (RAID 0) logical drives in the same array will also be failed.
- RAID 1 and RAID 1+0 configurations can tolerate multiple drive failures, as long as none of the failed drives are mirrored to one another.
- RAID 5 configurations can tolerate one drive failure.
- RAID ADG configurations can tolerate simultaneous failure of two drives in the array.
- Spare hard drive configuration guidelines include:
  - A spare must be assigned to each individual array separately.
  - A spare must be greater than or equal to any drive that it is intended to replace.
  - The same spare can be assigned to multiple Controllers as long as its capacity is greater than or equal to any drives in that array.

Note: Only remove hard drives that have been failed or marked as degraded by the Controller.

If more hard drives are failed than the fault-tolerance method allows, fault tolerance is “compromised” and the logical drive is failed. In this case, all requests from the operating system will be rejected with “unrecoverable” errors. The section on Compromised Fault Tolerance later in this chapter discusses possible ways to recover from this situation.

Recognizing Drive Failure

The indicators on the front of each hard drive are visible through the front of the external storage unit. When a drive is configured as a part of an array and attached to a powered-on Controller, the status of the drive can be determined from the illumination pattern of these indicators.

Figure 69 and its associated table define the three indicators on the hard drive.

For detailed descriptions of the various indicator combinations, refer to the “Hard Drives” section in the “Operation and Management” chapter.

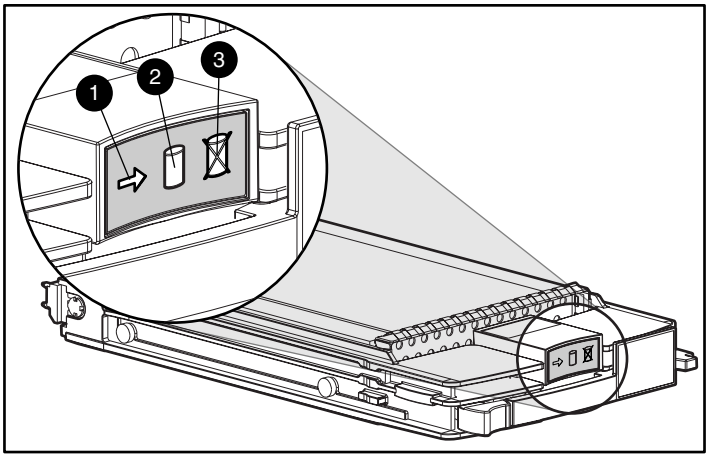

Figure 69: Hard drive indicators

Item	Description
❶	Environmental Monitoring Unit (EMU)
❷	Power
❸	Fault

There are several other ways to recognize that a hard drive has failed:

- The amber status indicator will light up on the front of an HP storage system if failed drives are inside. (Other problems such as fan failure, redundant power supply or over-temperature conditions will also cause this indicator to light up.)
- A front panel display message will list failed drives whenever the system is restarted, as long as the Controller detects one or more good drives.
- Compaq Insight Manager can detect failed drives locally or remotely across a network.

Additional information on troubleshooting hard drive problems can be found in the *Server Troubleshooting Guide*.

Compromised Fault Tolerance

Compromised fault tolerance commonly occurs when more physical drives have failed than the fault tolerance method can endure. In this case, the logical volume will be failed and error messages will be returned to the host. Data loss is likely to occur.

An example of this situation would be a RAID 5 logical drive, where one drive on an array fails while another drive in the same array is still being rebuilt.

Fault tolerance may also be compromised due to non-drive problems, such as a faulty cable, faulty storage system power supply, or a user accidentally turning off an external storage system while the host system power was on. In such cases, the physical drives do not need to be replaced. However, data loss may still have occurred, especially if the system was busy at the time the problem happened.

Procedure to Attempt Recovery

Inserting replacement drives when fault tolerance has been compromised will not improve the condition of the logical volume. Instead, if unrecoverable error messages display on your screen, try the following procedure to recover data.

1. Check for loose, dirty, broken, or bent cabling and connectors on all devices.
2. Turn the entire system off. Remove and then reinsert all hard drives and controllers.

Caution: Data can be lost if the drives are not firmly reseated.

3. Turn the system on. In some cases, a marginal drive might work again for long enough to allow you to make copies of important files.
4. If a #02 or #04 is displayed on the front panel, press the right push button to re-enable the logical volumes. Remember that data loss has probably occurred and any data on the logical volume is suspect.
5. Make copies of important data, if possible.
6. Replace any failed drives.
7. After the failed drives have been replaced, the fault tolerance may again be compromised. If so, cycle the power again, and if the #02 or #04 front-panel display message is displayed, press the right push button**.** This will re-enable the logical drive(s).

To minimize the risk of data loss from compromised fault tolerance, make frequent backups of all logical volumes.

Automatic Data Recovery

Automatic Data Recovery is an automatic background process that rebuilds data onto a spare or replacement drive when another drive in the array fails. The hard drive that is being rebuilt blinks once per second.

If a drive in a fault-tolerant configuration is replaced while the system power is off, a front-panel display message will be displayed during the next system startup. Automatic Data Recovery will be started.

When Automatic Data Recovery is completed, the online indicator of the replacement drive will stop blinking and begin to glow steadily.

In general, approximately 15 minutes is required to rebuild each gigabyte. The actual rebuild time depends upon:

- The rebuild priority level (high or low) of the logical drive
- The amount of I/O activity occurring during the rebuild operation
- The disk drive speed
- The number of drives in the array (for RAID 5 and RAID ADG)

For example, the rebuild time when using 9-GB Wide-Ultra hard drives in a RAID 5 configuration varies from 10 minutes per GB (for three drives) to 20 minutes per GB (for 14 drives).

Failure of Automatic Data Recovery

If the online indicator of the replacement drive stops blinking during Automatic Data Recovery, there are two possible causes:

- The replacement drive is failed (amber failure indicator illuminated or other indicators go out) and is producing unrecoverable disk errors. Remove and replace the failed replacement drive.
- The Automatic Data Recovery process may have abnormally terminated, due to a non-correctable read error (such as a SCSI bus signal integrity problem) from another physical drive during the recovery process.

Reboot the system and retry Automatic Data Recovery. If this does not work, backup all data on the system, do a surface analysis (using User Diagnostics), and restore the data from backup.

Replacing a Drive

Replacement drives must have a capacity no less than that of the smallest drive in the array. Drives of insufficient capacity will be failed immediately by the Controller before Automatic Data Recovery can begin.

Caution: Sometimes, a drive that has previously been failed by the Controller may seem to be operational after the system is power-cycled, or (for a hot-pluggable drive) after the drive has been removed and reinserted. However, continued use of such marginal drives may eventually result in data loss. Replace the marginal drive as soon as possible.

A hot-pluggable drive can be removed and replaced at any time, whether the host or storage system power is on or off. System performance and fault tolerance will both be affected until the rebuild operation is completed. This can take several hours, even if the system is not busy while the rebuild is in progress.

When a hot-pluggable drive is inserted, all disk activity on the array is paused while the new drive is spinning up (usually 10 seconds or so). If the drive is inserted while power is on, in a fault-tolerant configuration, data recovery onto the replacement drive will begin automatically (indicated by the blinking online LED).

Disk drive replacement should be performed as soon as possible. In addition, all logical volumes on the same array as the drive being replaced should have a current valid backup.

If another drive in the array should fail while you are rebuilding a previously failed drive, then fault tolerance has been compromised and all data on the array will be lost, with the following exceptions:

- In a mirrored (RAID 1) configuration, failure of a drive that is not mirrored to any other failed drives.
- A second drive failure in a RAID ADG configuration.

Take these precautions when removing failed drives to minimize the likelihood of compromising your logical drive data:

- Do not remove a degraded drive if any other member of the array is off-line (the online indicator is off).
- Do not remove a failed second drive from an array until the first failed or missing drive has been replaced *and* the rebuild process is complete. (When the rebuild is complete, the online indicator on the front of the drive will stop blinking.) Exceptions include RAID ADG configurations, as any two drives in the array can be replaced simultaneously.

When a hard drive is replaced, the Controller uses fault tolerance data on the remaining drives in the array to rebuild data (previously on the failed drive) onto the replacement drive. If more than one drive is removed at a time, the fault tolerance data is incomplete. The missing data cannot then be reconstructed and is likely to be permanently lost.

Moving Array Controller Drives and Arrays

Caution: Back up all data before removing drives or changing configurations. Failure to do so could result in permanent loss of data. Before moving drives and arrays, run the Array Configuration Utility.

Drives can be moved to other ID positions on the same array Controller. You may also move a complete array from one Controller to another, even if the Controllers are on different servers. Arrays on different Controllers can also be moved to another Controller.

Before moving drives, these conditions must be met:

- No drives are failed, missing, or degraded.
- The move will not result in more than 14 physical drives.
- No more than 32 logical volumes will be configured for a Controller.
- The array should be in its original configuration with no active spare drives.
- Capacity expansion is not running.
- Controller firmware is the latest version (recommended).

Before moving an array from one controller to another, these additional conditions must be met:

- All drives in the array must be moved at the same time.
- Positions of drives on the destination controller must not be changed during relocation of the array.

When the appropriate conditions are met, follow this procedure:

1. Power OFF the system.
2. Move the drives.
3. Power the system ON.

A #86 front-panel display message should be displayed, indicating that drive positions were changed and the configuration was updated. If a #121 front-panel display (no volumes) message is displayed, turn the system OFF immediately to avoid data loss, and return the drives to their original locations.

You can now check the new drive configuration by running the Array Configuration Utility.

Expanding and Extending Capacity

Caution: Back up all data before removing drives or changing configurations. Failure to do so could result in permanent loss of data. Before moving drives and arrays, run the Array Configuration Utility.

Array Capacity Expansion is the addition of physical drives to an array that has already been configured. The capacity of these added physical drives may then be added to an existing logical drive on the array (capacity *extension*; see the next paragraph), or it may be configured into a new logical drive.

Logical Drive Capacity Extension is the enlargement of an existing logical drive after the corresponding array has undergone capacity expansion.

Capacity expansion and extension are carried out using ACU. A data backup and restoration cycle is not required, even in non-fault-tolerant configurations.

If you are using hot-pluggable drives, expansion can be performed online (that is, without shutting down the operating system). Online *extension* can only be performed if supported by the operating system.

Note: When extending a logical drive under Windows 2000, upgrade the disk to DYNAMIC *before* creating a partition on that disk. If the disk already has a partition when it is upgraded to DYNAMIC, Windows 2000 may not allow logical drive extension. Refer to the Windows 2000 documentation for details about DYNAMIC and BASIC disks.

Note: Windows NT 4.0 will only allow four partitions on each logical drive. Additional drive space created may not be accessible if the four partitions per logical drive limit is exceeded.

Note: If you are running Windows 2000 with Microsoft Cluster Services (MSCS), logical drive extensions are not recommended. MSCS requires that disks be configured as BASIC in logical disk manager. In order to take advantage of logical drive extension, your drives would have to be configured as DYNAMIC when the volume is initially created. Because of the differences in requirements for the MSCS and the Logical drive extension feature, we recommend you do not perform a logical drive extension on a storage enclosure that is part of a Microsoft Cluster.

The expansion process is illustrated in the figure below where the original array (containing data) is shown with a dashed border, and the newly added drives are shown unshaded (they contain no data). The array controller redistributes the original logical drive over the enlarged array, using the same fault-tolerance method. The unused capacity on the new (enlarged) array can then be used to create an additional logical drive, with a different fault tolerance setting if necessary. Alternatively, the extra capacity can be used to increase the size of the original logical drive (capacity extension).

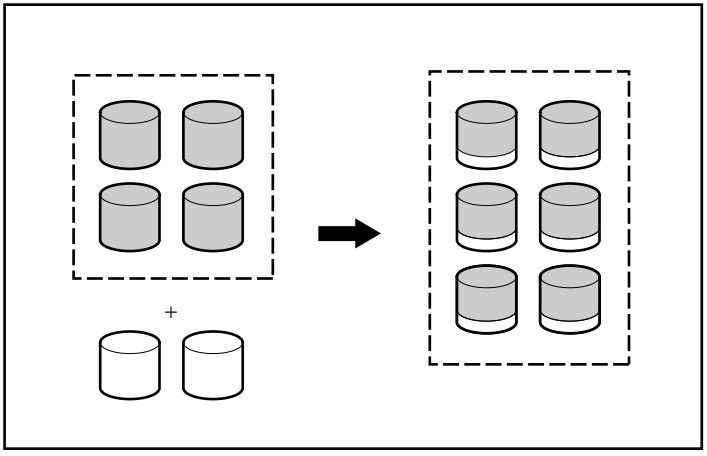

Figure 70: Array capacity expansion

If you are expanding an array that has several logical drives, data will be redistributed one logical drive at a time. Newly created logical drives will not become available until capacity expansion has finished.

Controller Display Messages

F

Each array controller in a MSA1000 contains an integrated display. This module is used for displaying informational and error messages, showing the current status of the module, and for providing user input when required. Traditional Power-On-Self-Test (POST) messages issued by PCI-based array controllers have been combined with runtime event notification messages to create a new set of controller display messages.

The display module consists of the following components:

- A two line, twenty column display text display window
- Four push buttons arranged in a circular “pie” shape
- Two status indicator lights

About the LCD Messages

The display module is capable of holding up to 100 messages. Once this maximum size has been reached, older messages will be removed to make room for newer ones. Messages can be of three types: error, informational, and user input.

The display message may specify a box number. The following box numbers are defined.

- Box 1 is the MSA1000 chassis.
- Box 2 is the storage enclosure attached to SCSI port A of the MSA1000.
- Box 3 is the storage enclosure attached to SCSI port B of the MSA1000.

Error Messages

Error messages indicate that a problem has occurred and may require user action to correct it. A complete list of possible messages and their meanings is contained in this appendix.

An amber indicator to the left of the display text display window will be turned on when ever an error message is currently being viewed. This indicator will also be turned on if an error message was sent to the display module but has not been viewed because non-error type messages were sent to the display module afterwards. By scrolling backwards and viewing all error messages, the indicator will return to only lighting up when currently on an error message.

Informational Messages

Informational messages indicate non-critical changes in the system that are provided as feedback to the user. A complete list of possible messages and their meanings is contained in this appendix.

The amber indicator to the left of the display text display window will normally be turned off whenever an informational message is currently being viewed unless an unviewed error message was previously sent to the display module. By scrolling backwards and viewing all error messages, the indicator will return to only lighting up when currently on an error message.

User Input Messages

User input messages indicate that the system has encountered an issue that can be handled in two different ways. It allows the user to either choose which way to handle the issue or it will default to a standard choice after a set period of time.

These user input messages will only occur during system power on and not during run time. A complete list of possible messages and their meanings is contained in this appendix.

The amber indicator to the left of the display text display window will blink on and off when a user input message is currently being viewed and is available for input. If the user has not provided input within the time-out period, the message will remain but the indicator will stop blinking.

Scrolling

Older messages can be viewed by scrolling backwards using the *up* push button (with the up arrow on it). Messages that are more recent can be viewed by scrolling forwards using the *down* push button (with the down arrow on it). The last message can be viewed by pressing the *left* button. When a new message is sent to the LCD, the display will show that message and ignore any previous scrolling position. This new message is now the most recent message available.

User Input

User input messages will allow the user two options, which are selected by either pushing the *left* push button (with the left arrow on it) or the *right* push button (with the right arrow on it). The meaning of each button will be defined in the display text display window.

Deleting Messages

The currently displayed message can be deleted from the display module by pressing the *left* push button and the *right* push button at the same time.

Redundancy Link Light

There is a green indicator to the left of the display text display window that will be lit when two array controllers have been inserted into the MSA1000 and they have enabled controller redundancy. The indicator will not be lit if there is only one array controller inserted or if the array controllers are not redundant due to some type of failure.

Note: You must have redundant cables connected to enable redundancy.

LCD Message Descriptions

The following table contains the defined messages and their components.

Table 20: LCD Message Descriptions

Message	Type	Description	Action
00 ARRAY CONTROLLER FIRMWARE VER <version>	Informational	Displays the current version of the firmware running on the array controller.	
01 MSA1000 STARTUP COMPLETE	Informational	The array controller has completed its power on sequence and is now operational.	
02 ENABLE VOLUME <n>? '<'=NO, '>'=YES	User Input	An issue has been found with a configured volume that may result in data loss. The exact nature of the issue will be detailed in a previous display message.	Selecting the *no* option will result in the volume being disabled so the user can attempt to fix the issue. Selecting the *yes* option will result in the volume being enabled regardless of the issue.
03 CRITICAL LOCK-UP DETECTED. CODE=<n>h	Error	A critical error has been detected by the array controller firmware. In order to prevent any possible data loss, the firmware has entered a lock-up state. The code contains engineering specific information about the lock-up condition. HP support should be contacted.	Remove the failing array controller, wait 10 seconds, and then reinsert it insuring that it is fully seated in the chassis. Should the issue persist please contact HP support.
04 ENABLE VOLUMES ? '<'=NO, '>'=YES	User Input	An issue has been found with all of the configured volumes that may result in data loss. The exact nature of the issue will be detailed in a previous display message.	An issue has been found with all of the configured volumes that may result in data loss. The exact nature of the issue will be detailed in a previous display message.

Table 20: LCD Message Descriptions

Message	Type	Description	Action
05 SYSTEM NAME: <name>	Informational	Displays the user assigned name for the MSA1000 system. This name can be assigned using the Array Configuration Utility (ACU).	
06 RESTARTING SYSTEM	Informational	Indicates that the system has been reset and is being restarted.	
20 INITIALIZING SCSI SUBSYSTEM	Informational	The SCSI subsystem is being initialized as part of the power on sequence.	
21 SCANNING FOR SCSI DEVICES	Informational	The firmware is searching for SCSI devices attached to the system as part of the power on sequence.	
22 INITIALIZING SCSI DEVICES	Informational	The firmware is initializing all SCSI devices attached to the system as part of the power on sequence.	
23 SCSI SUBSYSTEM HARDWARE FAILURE	Error	The SCSI subsystem on MSA1000 has experienced a hardware failure and is not operating correctly. The array controller has halted itself and cannot continue.	Please contact HP support.
24 BAD SCSI BUS MODE NON-LVD DEVICE FOUND	Error	The MSA1000 does not support SCSI Single Ended (SE) devices, it only supports SCSI Low Voltage Differential (LVD) devices.	The MSA1000 should be powered off and then all SCSI devices attached to it should be examined. Any SE devices found should be removed and replaced with LVD devices.

Table 20: LCD Message Descriptions

Message	Type	Description	Action
30 I2C READ FAILURE <I2C device name>	Error	The MSA1000 has a number of internal devices that are accessed via an I2C hardware bus. One of these devices failed when attempting to read from it. Certain I2C devices are considered critical and will result in a failure of the array controller while others may result in some loss of functionality (such as lost display messages).	Should the issue persist please contact HP support.
31 I2C WRITE FAILURE <I2C device name>	Error	The MSA1000 has a number of internal devices that are accessed via an I2C hardware bus. One of these devices failed when attempting to write to it. Certain I2C devices are considered critical and will result in a failure of the array controller while others may result in some loss of functionality (such as lost display messages).	Should the issue persist please contact HP support.
32 CHASSIS NVRAM CONTENTS CORRUPTED	Error	The MSA1000 has non-volatile memory on it that contains required information which is needed to operate. This non-volatile memory appears to be corrupted and the information is not valid. The MSA1000 cannot continue to operate and will halt.	Please contact HP support.
40 BEGIN REDUNDANCY SUPPORT	Informational	The array controllers are attempting to enter redundant mode.	

Table 20: LCD Message Descriptions

Message	Type	Description	Action
41 REDUNDANCY ACTIVE ACTIVE CONTROLLER	Informational	The array controllers are now in redundant mode and this array controller is *active* which means that it is allowed to access the configured volumes on the MSA1000.	
42 REDUNDANCY ACTIVE STANDBY CONTROLLER	Informational	The array controllers are now in redundant mode and this array controller is *standby* which means that it can be made *active* should the current *active* array controller fail, assuming you have all cables, I/O module or embedded switch installed.	
43 REDUNDANCY FAILED HARDWARE FAILURE	Error	While either attempting to enter redundant mode or already operating in redundant mode, one of the array controllers encountered a hardware failure on the communication channel between the two array controllers. Redundancy is disabled at this time.	If the MSA1000 is currently involved in host I/O, remove the *standby* array controller, wait 10 seconds, and then reinsert it insuring that it is fully seated in the chassis. If this does not resolve the issue then wait until down-time is available. Power off the MSA1000, remove both array controllers and reinsert them insuring they are fully seated in the chassis. Should the issue persist please contact HP support.

Table 20: LCD Message Descriptions

Message	Type	Description	Action
44 REDUNDANCY FAILED MISMATCH HARDWARE	Error	Both array controllers must contain the same hardware for them to successfully enter redundant mode. The current array controllers do not contain the same hardware, possibly because one has an attached fibre daughter card and the other does not.	If the MSA1000 is currently involved in host I/O, remove the *standby* array controller, add or remove the fibre daughter card as needed, wait 10 seconds, and then reinsert it insuring that it is fully seated in the chassis. If this does not resolve the issue then wait until down-time is available. Power off the MSA1000, remove both array controllers, add or remove fibre daughter cards as needed on both, and reinsert them insuring they are fully seated in the chassis. Should the issue persist please contact HP support.
45 REDUNDANCY FAILED MISMATCH FIRMWARE	Error	Both array controllers must be running the same version of firmware for them to successfully enter redundant mode. A process called firmware cloning that attempts to make them both the same firmware level has failed.	Manually update the firmware on the older array controller.
46 REDUNDANCY HALTED EXPAND ACTIVE	Error	Both array controllers must be running the same version of firmware for them to successfully enter redundant mode. A process called firmware cloning that attempts to make them both the same firmware level has failed.	

Table 20: LCD Message Descriptions

Message	Type	Description	Action
47 REDUNDANCY FAILED CACHE SIZE MISMATCH	Error	Both array controllers must have the same size of cache memory for them to successfully enter redundant mode.	If the MSA1000 is currently involved in host I/O, remove the *standby* array controller, add or remove cache memory as needed, wait 10 seconds, and then reinsert it insuring that it is fully seated in the chassis. If this does not resolve the issue then wait until down-time is available. Power off the MSA1000, remove both array controllers, add or remove cache memory as needed on both, and reinsert them insuring they are fully seated in the chassis. Should the issue persist please contact HP support.
48 REDUNDANCY HALTED FIRMWARE CLONED	Informational	Both array controllers must be running the same version of firmware for them to successfully enter redundant mode. A process called firmware cloning has been successfully completed in order to make them both the same firmware level. The *standby* array controller will now be restarted automatically so they can attempt to achieve redundancy again.	

Table 20: LCD Message Descriptions

Message	Type	Description	Action
49 REDUNDANCY FAILED FIRMWARE LOCKUP	Error	While either attempting to enter redundant mode or already operating in redundant mode, one of the array controllers encountered a critical condition resulting in a firmware lockup. Redundancy is disabled at this time.	If the MSA1000 is currently involved in host I/O, remove the *standby* array controller, wait 10 seconds, and then reinsert it insuring that it is fully seated in the chassis. If this does not resolve the issue then wait until down-time is available. Power off the MSA1000, remove both array controllers, and reinsert them insuring they are fully seated in the chassis. Should the issue persist please contact HP support.
50 REDUNDANCY FAILED OUT OF MEMORY	Error	While either attempting to enter redundant mode or already operating in redundant mode, one of the array controllers failed to allocate required memory. Redundancy is disabled at this time.	If the MSA1000 is currently involved in host I/O, remove the *standby* array controller, wait 10 seconds, and then reinsert it insuring that it is fully seated in the chassis. If this does not resolve the issue then you wait until down-time is available. Power off the MSA1000, remove both array controllers, and reinsert them insuring they are fully seated in the chassis. Should the issue persist please contact HP support.

Table 20: LCD Message Descriptions

Message	Type	Description	Action
51 REDUNDANCY FAILED I/O REQUEST ERROR	Error	While either attempting to enter redundant mode or already operating in redundant mode, one of the array controllers encountered an error while sending I/O between the two array controllers over the communication channel between them. Redundancy is disabled at this time.	If the MSA1000 is currently involved in host I/O, remove the *standby* array controller, wait 10 seconds, and then reinsert it insuring that it is fully seated in the chassis. If this does not resolve the issue then wait until down-time is available. Power off the MSA1000, remove both array controllers, and reinsert them insuring they are fully seated in the chassis. Should the issue persist please contact HP support.
52 REDUNDANCY FAILED PCI BUS ERROR	Error	While either attempting to enter redundant mode or already operating in redundant mode, one of the array controllers encountered a PCI bus error on the communication channel used between the two array controllers. Redundancy is disabled at this time.	If the MSA1000 is currently involved in host I/O, remove the *standby* array controller, wait 10 seconds, and then reinsert it insuring that it is fully seated in the chassis. If this does not resolve the issue then wait until down-time is available. Power off the MSA1000, remove both array controllers, and reinsert them insuring they are fully seated in the chassis. Should the issue persist please contact HP support.
53 REDUNDANCY FAILED NO SECOND CONTROLLER	Error	While operating in redundant mode, one of the array controllers was removed. Redundancy is disabled at this time.	Reinsert the missing array controller insuring that it is fully seated in the chassis.

Table 20: LCD Message Descriptions

Message	Type	Description	Action
54 REDUNDANCY FAILED CACHE DIMMS MISMATCH	Error	The cache memory modules on two different controllers are not the same size. All cache memory modules must be the same size for redundancy to operate.	Remove the array controller that has been halted, replace the cache modules with the appropriately sized ones, wait 10 seconds, and then reinsert the array controller insuring that it is fully seated in the chassis.
60 NO CACHE MODULE FOUND	Error	The array controller requires at least one cache module in order to operate. Either there is not one present or it has failed.	Remove the failed array controller, either add a cache module or replace the failed one, wait 10 seconds, and then reinsert it insuring that it is fully seated in the chassis. Should the issue persist please contact HP support.
61 DUAL CACHE MODULE SIZE MISMATCH	Error	The array controller has two cache modules attached but they are of different sizes. Both cache modules must be the same size.	Remove the failed array controller, replace one of the cache modules with a different one that is of the correct size, wait 10 seconds, and then reinsert it, insuring that it is fully seated in the chassis.
62 CACHE MODULE #<n> <n>MB	Informational	Displays the size of the cache module inserted into the respective cache module slot.	
63 VALID CACHE DATA FOUND AT POWER-UP	Informational	Valid host data was found in the battery backed cache memory at power up. This data has been flushed to the drives.	
64 CACHE DATA LOST BATTERY DEAD	Error	The battery on the cache memory was no longer charged. If there was data in the cache memory then it has been lost.	

Table 20: LCD Message Descriptions

Message	Type	Description	Action
65 CACHE HARDWARE ENABLED	Informational	The cache hardware had been temporarily disabled but is now enabled again. This may have been due to insufficient charge on the batteries that have now charged up to capacity.	
66 CACHE HARDWARE FAILED AND DISABLED	Error	The cache memory has experienced a hardware failure.	If the failure has occurred on the *standby* array controller, then remove the *standby* array controller, replace the cache modules, wait 10 seconds, and then reinsert the array controller insuring that it is fully seated in the chassis. If the failure has occurred on the *active* array controller, then wait until downtime is available. Power off the MSA1000, remove the array controller, replace the cache modules, and reinsert the array controller insuring it is fully seated in the chassis. Should the issue persist please contact HP support
67 CACHE HARDWARE TEMPORARILY DISABLED	Informational	The cache memory hardware has temporarily been disabled typically because either the battery is not charged up or a capacity expansion operation is occurring. The cache will automatically be enabled once the condition has been corrected.	

Table 20: LCD Message Descriptions

Message	Type	Description	Action
68 OBSOLETE CACHE DATA DELETED	Informational	Old data that no longer belongs to any current configured volumes was found in the cache memory at power up. This data has been deleted. This typically happens if cache modules are moved between array controllers.	
69 CACHE BATTERIES LOW, RECHARGING	Informational	The batteries on the cache module are low and are being recharged.	
70 CACHE DISABLED NO CONFIGURATION	Informational	The cache has not been configured and therefore is disabled. The cache can be configured by using the Array Configuration Utility (ACU).	

Table 20: LCD Message Descriptions

Message	Type	Description	Action
71 SYSTEM HALTED FOR CACHE ERROR	Error	This message is generated if the user chooses to ignore a critical cache error condition. It is always preceded by message #72 (although message #72 is removed from the LCD display once it has accepted user input).	The *no* option will cause the array controller to halt itself, providing the user an opportunity to resolve the issue. Selecting the *yes* option will result in the cache data being erased. Operation of the array controller will continue normally. **Error *1.1* and *1.2***: There is only one cache board in the array controller at present but it was previously configured with a second cache board that is now missing (dual cache module configuration). **Error *2.1* and *2.2***: A second cache board that contained valid data was removed from its original array controller and added to this array controller (dual cache module configuration). **Error *2.3***: A cache board that contained valid data was removed from its original array controller and added to this array controller (single cache module configuration). Return all cache boards to their original array controllers. Power up the systems without allowing any host I/O and wait for the cache data to be written to the drives. This will take a few minutes after the systems have finished the power on sequence. The systems can be powered off and cache boards moved to their new locations.

Table 20: LCD Message Descriptions

Message	Type	Description	Action
72 CACHE ERROR <n> IGNORE? <=NO >=YES	User Input	During power up, data was found in the cache that could not be flushed to the drives. The reason is either because the data does not belong to this array controller (the cache board was moved from a different array controller) or the cache data is partial (the rest of the data is in another cache board that was removed from the array controller). This error could occur if cache boards are moved improperly.	The *no* option will cause the array controller to halt itself, providing the user an opportunity to resolve the issue. Selecting the *yes* option will result in the cache data being erased. Operation of the array controller will continue normally. **Error *1.1* and *1.2***: There is only one cache board in the array controller at present but it was previously configured with a second cache board that is now missing (dual cache module configuration). **Error *2.1* and *2.2***: A second cache board that contained valid data was removed from its original array controller and added to this array controller (dual cache module configuration). **Error *2.3***: A cache board that contained valid data was removed from its original array controller and added to this array controller (single cache module configuration). Return all cache boards to their original array controllers. Power up the systems without allowing any host I/O and wait for the cache data to be written to the drives. This will take a few minutes after the systems have finished the power on sequence. The systems can be powered off and cache boards moved to their new locations.

Table 20: LCD Message Descriptions

Message	Type	Description	Action
73 CACHE HARDWARE BATTERIES MISSING	Error	The cache memory does not have its required batteries attached to it.	If the failure has occurred on the *standby* array controller, then remove the *standby* array controller, replace the cache modules, wait 10 seconds, and then reinsert the array controller insuring that it is fully seated in the chassis. If the failure has occurred on the *active* array controller, then wait until down-time is available. Power off the MSA1000, remove the array controller, replace the cache modules, and reinsert the array controller insuring it is fully seated in the chassis. Should the issue persist please contact HP support.
80 REPLACEMENT DRIVE FOUND BOX #<n> BAY <n>	Informational	A SCSI drive that was previously missing or failed has now been replaced with a working SCSI drive.	
81 SMART DRIVE ALERT BOX #<n>, BAY <n>	Informational	A SCSI drive may be close to failing. This was determined either by the drive firmware itself using SMART technology or by the array controller using monitor and performance testing.	The drive should be replaced as soon as possible following the guidelines in Appendix E, "Recovering from Hard Drive Failure" in this guide.
82 DRIVE HOT ADDED BOX #<n>, BAY <n>	Informational	A SCSI drive has been added to the MSA1000 or one of the storage enclosures attached to it.	

Table 20: LCD Message Descriptions

Message	Type	Description	Action
83 DRIVE HOT REMOVED BOX #<n>, BAY <n>	Informational	A SCSI drive has been removed from the MSA1000 or one of the storage enclosures attached to it.	
84 DRIVE FAILURE BOX #<n>, BAY <n> 84	Error	A SCSI drive in the MSA1000 or one of the storage enclosures attached to it has failed. If the drive was part of a configured volume, then the state of the volume will depend on the fault tolerance used.	The drive should be replaced as soon as possible following the guidelines in the "Recovering from Hard Drive Failure" section in this guide (Appendix E).
85 BAD DRIVE FRMWARE BOX #<n>, BAY <n>	Error	A SCSI drive has been detected that has known, bad firmware on it. Continued usage of this drive could result in drive failure, decreased performance or data loss.	Either the drive firmware should be updated or the drive should be replaced as soon as possible following the guidelines in the "Recovering from Hard Drive Failure" (Appendix E) in this guide.
86 DRIVE POSITION CHANGE DETECTED	Informational	The SCSI drives that make up a configured volume have been physically moved within the MSA1000 or an attached storage enclosure. The array controller has updated its configuration information accordingly.	
87 DRIVE POSITION CHANGE INVALID	Informational	The SCSI drives that make up a configured volume have been physically moved in such a way that the array controller can no longer access the configured volume.	The MSA1000 should be powered off and the drives restored to their original positions.

Table 20: LCD Message Descriptions

Message	Type	Description	Action
100 VOLUME #<n> STATE OK	Informational	The configured volume has returned to its normal operating state. This typically occurs after a rebuild operation has completed.	
101 VOLUME #<n> STATE FAILED	Error	The configured volume has been failed because too many SCSI drives that it is composed of have failed exceeding the fault tolerance level. The data on the configured volume is no longer available.	
102 VOLUME #<n> STATE INTERIM RECOVERY	Informational	The array controller has failed one or more SCSI drives that the configured volume is composed of but no data loss has occurred because fault tolerance is allowing the data to be recovered.	The failed drives should be replaced as soon as possible following the guidelines in the *Recovering from Hard Drive Failure* (Appendix E) section in this guide.
103 VOLUME #<n> STATE REBUILDING	Informational	The configured volume is rebuilding data on a SCSI drive that replaced a previously failed drive.	
104 VOLUME #<n> STATE DISABLED	Error	The configured volume has been disabled because too many of the SCSI drives that it is composed of are missing.	Power off the MSA1000 and then all attached storage enclosures. Unplug and reinsert all SCSI drives insuring they are fully seated in their bays. Check the cables connecting the MSA1000 to any attached storage enclosures. Power on the attached storage enclosures and then the MSA1000.
105 VOLUME #<n> STATE EXPANSION ACTIVE	Informational	The configured volume is currently performing a volume expansion operation.	

Table 20: LCD Message Descriptions

Message	Type	Description	Action
106 VOLUME #<n> STATE WAITING TO REBUILD	Informational	The configured volume is waiting to start rebuilding data on a SCSI drive that replaces a previously failed drive. The rebuild may not have started yet because the array controller is already performing a rebuild on another configured volume.	
107 VOLUME #<n> STATE WAITING TO EXPAND	Informational	The configured volume is waiting to start a volume expansion operation. The expansion may have not started yet because another configured volume is undergoing expansion or a rebuild is occurring on the configured volume.	
108 VOLUME #<n> STATE MISSING DRIVES	Error	The configured volume is missing too many of the SCSI drives that it is composed of making it unusable. The volume will be disabled.	Power off the MSA1000 and then all attached storage enclosures. Unplug and reinsert all SCSI drives insuring they are fully seated in their bays. Check the cables connecting the MSA1000 to any attached storage enclosures. Power on the attached storage enclosures and then the MSA1000.
109 VOLUME #<n> STATE WRONG DRIVE REPLACED	Error	The configured volume appears to have had known, good SCSI drives replaced instead of known, failed drives.	The MSA1000 should be powered off and the good drives should be restored while the failed drives should be replaced.

Table 20: LCD Message Descriptions

Message	Type	Description	Action
110 VOLUME #<n> EXPANSION DISABLED	Informational	The volume expansion operation on the configured volume has been disabled. This may be because a rebuild operation is ongoing, another expansion is already running, or the cache memory is disabled due to a low battery. The expansion will start once the condition has been cleared.	
111 VOLUME #<n> INITIALIZING PARITY	Informational	The array controller is calculating and storing parity information for the configured volume and therefore performance may be lower until it completes.	
112 VOLUME #<n> REBUILD FAILURE	Error	The rebuild operation on the configured volume has failed.	If the volume is still operating in regenerative mode, remove the new SCSI drive that was added as a replacement for the original failed drive and replace it with a different new drive.
113 VOLUME #<n> EXPANSION FAILURE	Error	The volume expansion operation on the configured volume has failed.	Run the Array Configuration Utility (ACU) and use it to determine the state of the volume. If the volume is still operational then it is possible to reattempt the operation.
114 VOLUME #<n> STATE DELETED	Informational	The configured volume has been deleted and is no longer available. Volumes are deleted by using the Array Configuration Utility (ACU).	
120 CONFIGURED VOLUMES <n>	Informational	The specified number of configured volumes were detected at power up.	

Table 20: LCD Message Descriptions

Message	Type	Description	Action
`121 NO VOLUMES DETECTED`	Informational	No configured volumes were detected at power up.	If there are supposed to be configured volumes, power off the MSA1000 and then all attached storage enclosures. Unplug and reinsert all SCSI drives insuring they are fully seated in their bays. Check the cables connecting the MSA1000 to any attached storage enclosures. Power on the attached storage enclosures and then the MSA1000.
`122 NEW VOLUME(S) DETECTED`	Informational	Configured volumes from another array controller were migrated to this array controller. The configuration information has been updated.	
`123 TOO MANY VOLUMES DETECTED`	Error	The array controller only supports a maximum of 32 configured volumes. More volumes then that were detected at power up. This typically occurs when migrating a set of volumes from one array controller to a different array controller that already has configured volumes on it. The migrated volumes have not been added.	Remove the migrated drives and run the Array Configuration Utility (ACU). Delete any unneeded volumes until the number of existing volumes plus the number of migrated volumes is 32 or less. Add the migrated drives back.

Table 20: LCD Message Descriptions

Message	Type	Description	Action
125 ACCESS CONTROL CONFLICT DETECTED	Error	A set of volumes have been migrated from one array controller to a different array controller that already has configured volumes on it. The migrated volumes have access controls defined for them that conflicts with the existing configuration. The access controls has been modified so as to allow the migration to proceed.	Run the Array Configuration Utility (ACU) to check the new access controls and modify them if needed.
126 ACCESS CONTROL RESOURCES EXCEEDED	Error	A set of volumes have been migrated from one array controller to a different array controllers that already has configured volumes on it. The migrated volumes have access controls defined for them that conflicts with the existing configuration. The access controls has been modified so as to allow the migration to proceed.	Run the Array Configuration Utility (ACU) to check the new access controls and modify them if needed.
201 ARRAY CONTROLLER TEMPERATURE OK	Informational	The temperature sensor on the array controller indicates that the temperature which was previously exceeding the normal operating range is now back within the range.	

Table 20: LCD Message Descriptions

Message	Type	Description	Action
202 ARRAY CONTROLLER OVERHEATING	Error	The temperature sensor on the array controller indicates that the array controller is starting to exceed the normal operating range.	Check all MSA1000 fans and insure they are operating. Any failed fans should be replaced. Insure that there are drive blank cartridges in any empty drive bays of the MSA1000 chassis. If only one array controller is inserted, insure that there are cover plates installed in the empty array controller bay and the fibre bay of the MSA1000 chassis.
203 ARRAY CONTROLLER OVERHEATED	Error	The temperature sensor on the array controller indicates that the array controller has exceeded the safe operating range.	The MSA1000 should be powered off as soon as possible to avoid hardware failure. Check all MSA1000 fans and insure they are operating. Any failed fans should be replaced. Insure that there are drive blank cartridges in any empty drive bays of the MSA1000 chassis. If only one array controller is inserted, insure that there are cover plates installed in the empty array controller bay and the fibre bay of the MSA1000 chassis.
204 ARRAY CONTROLLER DISABLED	Error	The array controller has been disabled due to a redundancy failure.	Remove the failed array controller, wait 10 seconds, and then reinsert the array controller insuring that it is fully seated in the chassis. Should the issue persist please contact HP support.
205 ARRAY CONTROLLER RESTARTING	Informational	The array controller has completed firmware cloning and will be restarted automatically.	

Table 20: LCD Message Descriptions

Message	Type	Description	Action
300 RECOVERY ROM AUTOFLASH STARTED	Informational	Indicates that the array controller has detected that the firmware's backup recovery ROM image is invalid and is copying the current active firmware image into the backup recovery ROM.	
301 RECOVERY ROM AUTOFLASH DONE	Informational	Indicates that the array controller has successfully completed the process of copying the current active firmware image into the backup recovery ROM.	
302 RECOVERY ROM AUTOFLASH FAILED	Error	Indicates that the array controller failed to copy the current active firmware image into the backup recovery ROM. Recovery ROM support is disabled.	Remove the failing array controller, wait 10 seconds, and then reinsert the array controller insuring that it is fully seated in the chassis. The ROM autoflash process will be attempted again. Should the issue persist please contact HP support.

Table 20: LCD Message Descriptions

Message	Type	Description	Action
`303 ROM CLONING STARTED`	Informational	Indicates that the two array controllers in a MSA1000 do not have the same version of firmware on them. Therefore, one array controller's version of the firmware will be copied on to the other array controller. Both controllers must be running the same version of firmware in order for controller redundancy to operate. If both array controllers are in the power up sequence, then the most recent version of firmware will be used. If one array controller has already completed the power up sequence and is now *active*, then its version of firmware will be used even if less recent.	
`304 ROM CLONING DONE`	Informational	Indicates that the two array controllers in a MSA1000 have finished copying one array controller's version of firmware to the other array controller. Both controllers must be running the same version of firmware in order for controller redundancy to operate.	
`305 ROM CLONING FAILED`	Error	Indicates that the two array controllers in a MSA1000 failed to copy one array controller's version of firmware to the other array controller. Both controllers must be running the same version of firmware in order for controller redundancy to operate.	Remove the *standby* array controller, wait 10 seconds, and then reinsert the array controller insuring that it is fully seated in the chassis. The ROM cloning process will be attempted again. Should the issue persist please contact HP support.

Table 20: LCD Message Descriptions

Message	Type	Description	Action
306 FIRMWARE FLASH STARTED	Informational	Indicates that the array controller in the MSA1000 has started the firmware flash process. Do not turn off power to the system until it has completed. This could take several minutes.	
307 FIRMWARE FLASH DONE	Informational	Indicates that the array controller in the MSA1000 has completed the firmware flash process. It is now safe to turn off power to the system.	
308 FIRMWARE FLASH FAILED	Error	Indicates that the array controller in the MSA1000 has failed the firmware flash process.	Attempt the flash process again. Should the issue persist, please contact HP support.
309 EMU FLASH STARTED	Informational	Indicates that the EMU in the MSA1000 has started the firmware flash process. Do not turn off the power to the system until it has completed. This could take five minutes.	
310 EMU FLASH DONE	Informational	Indicates that the EMU in a MSA1000 has completed the firmware flash process. It is now safe to turn off power to the system.	
311 EMU FLASH FAILED	Error	Indicates that the EMU in a MSA1000 has failed the firmware flash process.	Attempt the flash process again. Should the issue persist, please contact HP support.
400 STORAGE BOX #<n> FAN OK	Informational	The specified storage enclosure indicates that one of its fans which previously had been failed or degraded is now operating normally.	

Table 20: LCD Message Descriptions

Message	Type	Description	Action
401 STORAGE BOX #<n> FAN FAILED	Error	The specified storage enclosure indicates that one of its fans has failed. The storage enclosure and any devices in it may now be susceptible to overheating if corrective action is not taken.	Check all fans and insure they are operating. Any failed fans should be replaced.
402 STORAGE BOX #<n> FAN DEGRADED	Error	The specified storage enclosure indicates that one of its fans is not operating at full efficiency. The fan may eventually fail.	Check all fans and insure they are operating. Any failed fans should be replaced.
403 STORAGE BOX #<n> FAN HOT INSERTED	Informational	The specified storage enclosure indicates that a fan has been added.	
404 STORAGE BOX #<n> FAN HOT REMOVED	Informational	The specified storage enclosure indicates that a fan has been removed.	
405 STORAGE BOX #<n> TEMPERATURE OK	Informational	The temperature sensor in the storage enclosure indicates that the temperature is now back in the normal operating range.	
406 STORAGE BOX #<n> OVERHEATING	Error	The temperature sensor in the storage enclosure indicates that the enclosure is starting to exceed the normal operating range.	Check all fans and insure they are operating. Any failed fans should be replaced. Insure that there are drive blank cartridges in any empty drive bays in the enclosure. If the enclosure is a MSA1000 and only one array controller is inserted, insure that there are cover plates installed in the empty array controller bay and the fibre bay in the chassis.

Table 20: LCD Message Descriptions

Message	Type	Description	Action
407 STORAGE BOX #<n> OVERHEATED	Error	The temperature sensor in the storage enclosure indicates that the enclosure has exceeded the safe operating range.	The MSA1000 should be powered off as soon as possible and immediately after that, the enclosure should be powered off, to avoid hardware failure. Check all fans and insure they are operating. Any failed fans should be replaced. Insure that there are drive blank cartridges in any empty drive bays in the enclosure. If the enclosure is a MSA1000 and only one array controller is inserted, insure that there are cover plates installed in the empty array controller bay and the fibre bay in the chassis.
408 STORAGE BOX #<n> POWER SUPPLY OK	Informational	The specified storage enclosure indicates that one of its power supplies which previously had been failed is now operating normally.	
409 STORAGE BOX #<n> POWER SUPPLY FAILED	Error	The specified storage enclosure indicates that one of its power supplies has failed.	Check all power supplies and insure they are operating. Any failed power supplies should be replaced.
410 STORAGE BOX #<n> POWER SUPPLY ADDED	Informational	The specified storage enclosure indicates that a power supply has been added.	
411 STORAGE BOX #<n> POWER SUPPLY REMOVED	Informational	The specified storage enclosure indicates that a power supply has been removed.	

Table 20: LCD Message Descriptions

Message	Type	Description	Action
412 STORAGE BOX #<n> EMU NOT RESPONDING	Error	The specified storage enclosure is not responding to commands.	Insure the storage box is powered on. Insure all cables are connected securely. Power of the MSA1000 and the storage box. Power on the storage box first and then the MSA1000. Should the issue persist, please contact HP support.
413 STORAGE BOX #<n> EMU VERSION <version>	Informational	The version of the firmware running on the EMU. This is only displayed for the internal EMU of the MSA1000. It is not displayed for externally connected storage boxes.	
415 STORAGE BOX #2 OR #3 HOT ADDED	Informational	A HP StorageWorks SCSI expansion storage enclosure has been hot-added to the MSA1000.	
500 INITIALIZING PCI SUBSYSTEM	Informational	The array controller's PCI subsystem is being initialized as part of the power up sequence.	
501 PCI SUBSYSTEM HARDWARE FAILURE	Error	The array controller's PCI subsystem has encountered a critical error during the power up sequence.	Remove the failed array controller, wait 10 seconds, and then reinsert the array controller insuring that it is fully seated in the chassis. Should the issue persist please contact HP support.
502 PCI BRIDGE ASIC SELF TEST FAILURE	Error	The array controller's PCI bridge ASIC has encountered a critical error during the power up sequence.	Remove the failed array controller, wait 10 seconds, and then reinsert the array controller insuring that it is fully seated in the chassis. Should the issue persist please contact HP support.

Table 20: LCD Message Descriptions

Message	Type	Description	Action
510 INITIALIZING FIBRE SUBSYSTEM	Informational	The array controller's fibre subsystem is being initialized as part of the power up sequence.	
513 UNCORRECTED ECC MEMORY ERROR SEEN	Error	The array controller has detected an uncorrectable error in the ECC memory on the memory cache board.	Remove the failed array controller and replace the memory cache board with a new one.
514 FIBRE SWITCH HARDWARE FAILURE	Error	The MSA SAN Switch 2/8 that is installed in the MSA1000 failed to establish a connection with the array controller over fibre channel.	Remove the MSA Fabric Switch 6 switch, wait one minute and then reinsert it insuring it is fully seated in the chassis. Wait one minute and check the LCD to see if a new copy of this error message is created. The service indicator LED on the back of the switch should turn solid green if the switch is operating normally. The service indicator LED will flash amber if it has failed again. Should the issue persist please contact HP support.

Table 20: LCD Message Descriptions

Message	Type	Description	Action
515 FIBRE DEVICE HARDWARE FAILURE	Error	The fibre device that is installed in the MSA1000 fibre bay was not recognized by the array controller.	Verify the fibre device is supported by the MSA1000 by checking the documentation that came with it. If it is supported then remove the fibre device, wait one minute and then reinsert it insuring it is fully seated in the chassis. Wait one minute and check the LCD to see if a new copy of this error message is created. The service indicator LED on the back of the switch should turn solid green if the device is operating normally. The service indicator LED will flash amber if it has failed again. Should the issue persist, please contact HP support.
516 FIBRE SUSBYSTEM LINK FAILURE	Error	There is no active fibre connection to this MSA1000 array controller (Laser OFF). If the Fibre connection is a direct-connect from an HBA to the MSA1000, this message is expected when the server is powered off or restarted.	Power on the server and load the HBA drivers. If the status does not change to OK, check cables, fibre bay board, SFP, and HBA.
517 FIBRE SUBSYSTEM LINK OK	Informational	There is an active fibre connection to this MSA1000 array controller (Laser On). This message is only displayed when preceded by message 516.	

Table 20: LCD Message Descriptions

Message	Type	Description	Action
518 PERSISTENT MEM ENABLED	Informational	Global variables such as system prompts and profile information remain persistent in cache over power cycles of the MSA1000. This message is displayed each time the MSA1000 is powered up.	

Recovery ROM and ROM Cloning

Recovery ROM

Each MSA1000 Controller contains ROM (Read-Only Memory), which holds the firmware that operates the controller. The Recovery ROM feature stores two complete firmware images in the ROM: one active image and one backup image. When the controller is powering up, it checks both firmware images to ensure they are valid. If either one is not, the valid image will be copied on top of the invalid image to correct it. This is referred to as auto-flashing. All of this functionality is done automatically by the controller and does not require any user intervention.

ROM Cloning

Note: In order for a MSA1000 to operate in a redundant controller configuration, it must contain two controllers that are executing the same version of firmware. During power up (or if an optional controller is hot-plugged while the MSA1000 is already operating) the firmware versions on both controllers are compared. If they are not the same, then the ROM Cloning feature will attempt to copy one version of firmware onto the other controller. After the copy has been completed, the controller that was modified will be automatically reset. Once the reset controller has powered up the two controllers should then start redundant operation. All of this functionality is done automatically by the controllers and does not require any user intervention.

The determination as to which firmware version will be used is based on the following criteria:

- If the MSA1000 is being powered up with both controllers inserted, then the most recent version of firmware will be used regardless of which controller it resides.
- If the MSA1000 is already operating and an optional controller is hot-plugged, then the version of firmware that is on the original (non-hot-plugged) controller will be used regardless of its version. This ensures that any host initiated I/O to the controller is not interrupted.

There is the possibility that a specific version of firmware may not be compatible with certain hardware revisions of a controller. In this scenario, the most recent firmware version that is compatible with both controllers will be copied to the controller with the incompatible firmware version. However, if the controller that is updated is already operating and processing I/O, then it will not be reset. The MSA1000 will not enter redundant operation and an appropriate message will be shown on the display. After the MSA1000 has been shutdown and powered back on, the controllers will then be able to enter redundant operation. On a subsequent power cycle, both controllers will enter redundant mode.

SCSI ID Assignments

H

SCSI IDs are assigned automatically in the MSA1000 drive shelf and any additional attached storage enclosures according to the drive bay used for each drive.

The tables in this appendix show the assigned SCSI assignments and the SCSI buses used for each of the drive bays in the MSA1000 and any additional attached storage enclosures.

Note: Although the 4214 and the 4254 storage enclosures can be attached to the MSA1000, optimum performance is obtained when using 4314 and 4354 storage enclosures.

A separate table is present on each page of this appendix.

SCSI ID Assignments for the MSA1000

Table 21 shows the SCSI ID Assignments for the MSA1000 drive shelf.

Table 21: MSA1000 SCSI ID Assignments

Drive Bay	SCSI ID	SCSI Bus
1	0	0
2	1	0
3	2	0
4	3	0
5	4	0
6	5	0
7	8	0
8	0	1
9	1	1
10	2	1
11	3	1
12	4	1
13	5	1
14	8	1

SCSI ID Assignments for 4314 Storage Enclosures

Table 22 shows the SCSI assignments for additional 4314 storage enclosures attached to the MSA1000.

Table 22: Additional 4314 Storage Enclosure SCSI ID Assignments

Drive Bay of Each Enclosure	SCSI IDs	SCSI Bus, first Additional Enclosure	SCSI Bus, Second Additional Enclosure
1	0	2	3
2	1	2	3
3	2	2	3
4	3	2	3
5	4	2	3
6	5	2	3
7	8	2	3
8	9	2	3
9	10	2	3
10	11	2	3
11	12	2	3
12	13	2	3
13	14	2	3
14	15	2	3

SCSI ID Assignments for 4354 Storage Enclosures

Table 23 shows the SCSI assignments for an additional 4354 storage enclosure attached to the MSA1000.

Table 23: Additional 4354 Storage Enclosure

Drive Bay	SCSI ID	SCSI Bus
1	0	2
2	1	2
3	2	2
4	3	2
5	4	2
6	5	2
7	8	2
8	0	3
9	1	3
10	2	3
11	3	3
12	4	3
13	5	3
14	8	3

index

A

access control list 14
Access Control Lists See ACL
ACL
  adding to with the CLI 157
  deleting from, with the CLI 159
  disabling, with the CLI 162
  viewing, with the CLI 157
ACU
  accelerator read-write ratio 102
  accessing 85
  array
    creating 93, 103
    deleting 105
    expanding 106
  array accelerator, disabling 102
  automatic configuration, See express configuration mode
  cache read-write ratio 102
  clearing the configuration 101
  configuration
    clearing 101
    modes of 90
    modifying 100
    of a new controller 90
    of hardware 13
    screen types 88
  controller settings menu option 102
  creating
    array 93, 103
    logical drive 96, 104
deleting
  array 105
  logical drive 106
disabling the array accelerator 102
display settings 84
drive failure, probability of 114
enabling the array accelarator 102
expanding an array
  procedure for 106
  setting priority of 102
express configuration mode
  availability of 86
  for a new controller 91
  for a pre-configured controller 100
extending a logical drive 108
features, overview of 83
Insight Manager, using 87
installing 85
logical drive
  creating 96, 104
  deleting 106
  extending 108
  migrating 109
manual configuration
  of a new controller 93, 98
  of a pre-configured controller 112
methods of accessing 85
migrating a logical drive 109
modifying a configuration 100
priority settings 102
RAID level, migration of 109
read-write ratio 102
rebuild priority setting 102

reconfiguring a controller 101
removing a spare drive 110
screen resolution requirements 84
setting the expand or rebuild priority 102
Smart Component, using 85
spare drives
adding or removing 110
SSP (Selective Storage Presentation) 111
standard configuration mode
for a new controller 98
stripe size
migrating 109
time required for expansion, extension, or migration 107, 109
wizards, using
for a new controller 93
for a pre-configured controller 101
adding hard drive to array 192
ADG See RAID ADG
advanced data guarding (RAID ADG) 180
advanced data guarding See RAID ADG
amber LED 185
array 171
accelerator batteries, installing 67
adding hard drives 192
capacity expansion 192
creating 93, 103
defined 174
moving 191
online spares in 175
physical limitations of 174
array accelerator
described 11
features of 11
array configuration methods 13
Array Configuration Utility, see ACU 83
array controller configuration, with the CLI 134
array controller, duplexing 182
audience xiv
automatic data recovery
failure of 188
limitation of 182
resources 188

B

batteries
replacement notice 166
replacing 67
warnings 67
block of data, defined 173

C

cable connectors, illustrated 37
cables
customizing 37
managing 37
rack systems 37
troubleshooting 187
cache, enabling/disabling per LUN with the CLI 151
capacity expansion 192
capacity extension 192
chassis 62
clearing the configuration 101
CLI
ACL commands 152, 156
ACL, adding to 157
ACL, deleting from 159
ACL, disabling 162
ACL, viewing 157
array controller configuration commands 134
array, expanding 148
command option, defined 119
command syntax 119
connections, changing the HBA 161
connections, changing the name of 160
connections, deleting names of 161
connections, managing profiles of 155
connections, naming 154
connections, viewing 153, 157
controller ID, setting of 135
controller settings, viewing 130, 132, 133

disks, viewing information about 127
display commands 127
global settings 134
hard drives, locating 137
help command, described 124
Hyper Terminal, setting up 121
inter-controller link, defined 118
LUN management commands 137
LUN names, assigning 145
LUN names, viewing 130
LUN, adding a spare to 145
LUN, cache, enabling/disabling 151
LUN, creating 140
LUN, deleting 146
LUN, extending 149
LUN, RAID level, changing 150
LUN, spare, including 144
LUN, viewing information about 128
overview of 118
prompt, changing of 136
serial cable requirements 121
serial connection, setup 121
setup 121
spare drives, deleting 147
SSP commands 152, 156
command line interface, see CLI
Compaq Insight Manager
drive failure detection 186
indicating faults 52
comparison
of hardware-based RAID with software-based RAID 182
of logical drive failure risk for different RAID levels 115
of RAID methods with other fault-tolerance methods 182
components
hardware, of the MSA1000 9
hot-pluggable 7
replacing 7
software, of the MSA1000 20
compromised fault tolerance 186
configuration
fault-tolerant 13
of drives 13
configuration needs, specific
Linux 27
NetWare 28
OpenVMS 29
Tru64 UNIX 31
Windows 26
configurations, samples of 42
connections
changing the HBA, with the CLI 161
changing the name of, with the CLI 160
deleting names of, with the CLI 161
managing profiles of, with the CLI 155
naming, with the CLI 154
viewing, with the CLI 153, 157
Controller
removing 62
Controller display
box numbering 196
components 195
defined 10, 195
deleting messages 197
error messages 10, 196
illustrated 10
informational messages 196
message types 196
redundancy link light 197
user input messages 196, 197
Controller display messages
access control conflict detected 217
access control resources exceeded 217
array Controller disabled 218
array Controller overheated 218
array Controller overheating 218
array Controller restarting 218
array Controller temperature ok 217
bad drive firmware box, bay 212
bad SCSI bus mode non-lvd device found 199
begin redundancy support 200

cache batteries low recharging 208
cache data lost battery dead 206
cache disabled no configuration 208
cache error 210
cache hardware batteries missing 211
cache hardware enabled 207
cache hardware temporarily disabled 207
cache module size 206
chassis nvram contents corrupted 200
configured volumes 215
critical lockup detected 198
drive failure box, bay 212
drive hot added box, bay 211
drive hot removed box, bay 212
drive position change detected 212
drive position change invalid 212
dual cache module size mismatch 206
EMU flash done 221
EMU flash failed 221
EMU flash started 221
enable volume 198
enable volumes 198
fibre switch hardware failure 225
firmware flash done 221
firmware flash failed 221
firmware flash started 221
firmware version 198
initializing fibre subsystem 225
initializing PCI subsystem 224
initializing SCSI devices 199
initializing SCSI subsystem 199
new volume(s) detected 216
no cache module found 206
no volumes detected 216
obsolete cache data deleted 208
PCI bridge ASIC self-test failure 224
PCI subsystem hardware failure 224
read failure 200
recovery ROM autoflash done 219
recovery ROM autoflash failed 219
recovery ROM autoflash started 219
redundancy active active Controller 201
redundancy active standby Controller 201
redundancy failed cache DIMMS mismatch 206
redundancy failed cache size mismatch 203
redundancy failed firmware lockup 204
redundancy failed hardware failure 201
redundancy failed I/O request error 205
redundancy failed mismatch firmware 202
redundancy failed mismatch hardware 202
redundancy failed no second Controller 205
redundancy failed out of memory 204
redundancy failed PCI bus error 205
redundancy halted expand active 202
redundancy halted firmware cloned 203
replacement drive found box 211
restarting system 199
ROM cloning done 220
ROM cloning failed 220
ROM cloning started 220
scanning for SCSI devices 199
SCSI subsystem hardware failure 199
smart drive alert box 211
startup complete 198
storage box 221
storage box EMU not responding 224
storage box EMU version 224
storage box fan failed 222
storage box fan hot inserted 222
storage box fan hot removed 222
storage box overheated 223
storage box overheating 222
storage box power supply added 223
storage box power supply failed 223
storage box power supply ok 223
storage box power supply removed 223
storage box temperature ok 222
storage fan degraded 222
system halted for cache error 209
system name 199
too many volumes detected 216
uncorrected ECC memory error seen 225
valid cache data found at power-up 206

volume expansion disabled 215
volume expansion failure 215
volume initializing parity 215
volume rebuild failure 215
volume state deleted 215
volume state disabled 213
volume state expansion active 213
volume state failed 213
volume state ok 213
volume state rebuilding 213
volume state waiting to expand 214
volume state waiting to rebuild 214
volume state wrong drive replaced 214
volumestate interim recovery 213
volumestate missing drives 214
write failure 200
controller duplexing 182
controller settings, viewing, with the CLI 130, 132, 133
conventions
document xv
equipment symbols xvi
text symbols xvi

D

DAS See Direct Attached Storage 79
data
distributing 13
loss of 183
rebuild time 188
recovery,automatic 188
restoring
automatically 7
data block, defined 173
data protection methods
non-RAID 182
RAID 175
data stripes, defined 173
deployment overview 2
direct attached storage, capacity expansion 79
disks, viewing information about, with the CLI 127
distributed data guarding (RAID 5) 179
document conventions xv
documentation, related xiv
drive
arrays, performance 21
bays, IDs 8
failure, notification 186
failure, replacing drive 189
logical 21
performance 13
physical 13
drive array, See array
drive failure probability, graph of 115
drive mirroring (RAID 1+0) 177
duplexing, controller 182

E

electrostatic discharge 167
EMU
illustrated 15
indicators 56
enclosure expansion
ACU 80
adding enclosures to existing MSA1000 81
adding SW 4314/SW4214 enclosures 79
attaching SCSI cables 79, 81
direct attached storage 79
installation overview 81
methods of 79
migration overview 79
rack considerations 79
recommended procedures 79
running disk administration utility 80
SCSI connectors, illustrated 80
Ultra3 drives 79
enclosure status indicators 54
enclosure status indicators, illustrated 54
enlarging logical drive 192
equipment symbols xvi
expanding
array capacity 192

expanding an array
  procedure for 106
extending
  logical drive capacity 192
  logical drive, with the ACU 108

F

failed
  hard drive 184
  logical drive 184, 186
fault management
  automatic data recovery 188
fault tolerance
  alternative methods of 182
  compromised 186
  controller duplexing as 182
  definition 13
  description of methods 175
  online spare 13
  removing hard drives 75
  See also RAID methods
  software-based RAID as 182
FCC notice 163
features
  front view 5
  rear view 6
Fibre Channel Arrays, applying power 51
Fibre Channel I/O Module
  hot-pluggability 7, 18
  illustrated 17
  indicators 59
  redundant 17
  replacement 7, 17, 18, 19
Fibre Channel I/O module
  replacement 18
Fibre Channel I/O module indicators, illustrated 59
firmware
  updating 82
flashing LEDs, with the CLI 137
front view features 5
front view, MSA1000, illustrated 5

G

getting help xviii
global controller settings, changing with the CLI 134
graph, drive failure probability 115
grounding methods 168

H

hard drives
  adding to array 192
  configuring 51
  defining arrays 13
  ejector levers 76, 77
  failure 184
    detection 186
    multiple 183
    protection against 175
    recognizing 185
    replacing drive 189
  hot-pluggable 7
  identifying 51
  indicators 59
    amber 59
    illustrated 57
  latches 76, 77
  moving 191
  online spare 13
  recognizing 77
  reconstructing 77
  removing 76
  removing in fault-tolerant systems 75
  replacement 7
  replacing 189
    conditions 7
    procedure 76
  securing 77
  status lights 185
hard drives, locating with the CLI 137
hardware
  failure 13

HBA supported
Linux 27
NetWare 28
OpenVMS 29
Tru64 UNIX 31
Windows 26
HBA, supported 36
help, obtaining xviii
hot spare 175
hot-plug drive, replacing 189
hot-pluggability
definition 7
hp
authorized reseller xix
technical support xviii
website xix
Hyper Terminal, setting up 121

I

I/O
rates 13
indicators 5
amber 52
drive access 76
drive tray 59, 77
EMU 56
enclosure status 54
Fibre Channel I/O module 59
interpreting 185
online 76
power supply/blower assembly 55
using 52
visibility 37
Insight Manager
monitoring performance 21
installing
batteries 67
inter-controller link, defined 118
interim data recovery, limitation of 182

L

latches, hard drive 76, 77
limitations
moving arrays 191
moving drives 191
Linux environments
specific configuration needs 27
supported HBA 27
supported operating systems 26
supported software components 26
Linux environments, information about 26
load balancing, defined 177
logical drive
capacity extension 192
compared to array 174
creating 104
defined 172, 173
failure 184, 186
migrating 109
probability of failure 115
LUN
changing the RAID level with the CLI 150
creating with a spare, with the CLI 144
creating with the CLI 140
deleting with the CLI 146
extending with the CLI 149
names, assigning with the CLI 145
names, viewing, with the CLI 130
viewing information about, with the CLI 128

M

manual configuration
of a new controller 98
maximum storage 8
migrating a logical drive 109
mirroring of drives 177
modifying a configuration 100
moving
array 191
drives 191
MSA Hub 2/3

defined 19
hot-pluggability 7
illustrated 19
redundant 19
replacement 7
MSA SAN Switch
defined 18
full duplex non-blocking performance 18
hot-pluggability 7
illustrated 18
redundant 18
replacement 7
MSA1000 Controller
array accelerator 11
array accelerator features 11
ejector levers 62, 64
failure 62
hot-pluggability 7
indicators 52
installing 63
latches 62, 63, 64
managing hard drive array 13
removing 62
replacement 7
replacing 62
securing 63
MSA1000 startup sequence 51

N

NetWare environments
specific configuration needs 28
supported HBA 28
supported operating systems 27
supported software components 28
NetWare environments, information about 27
no fault tolerance (RAID 0) 176

O

online spare
defined 175
limitations of 175
use of, with alternative fault-tolerance methods 182
OpenVMS environments
specific configuration needs 29
supported HBA 29
supported operating systems 28
supported software components 29
OpenVMS environments, information about 28
operating systems
supported 23
operating systems supported
Linux 26
NetWare 27
OpenVMS 28
Tru64 UNIX 31
Windows 25
overview
of the MSA1000 2

P

panels
front 51
rear 63
parity data
in RAID 5 179
in RAID ADG 180
performance
logical drives 13
physical drives, See hard drives
power
applying 51
connecting 50
cord 50
switch 51
system 8
power cords
disconnecting 8
MSA1000 50
power supplies
adding 16
hot-pluggability 7
illustrated 16

redundant 7, 16
removing 37
removing power from 8
replacement 7
power supply/blower assembly indicators 55
power switch
definition of positions 8
location of 5
prerequisites
moving arrays 191
moving hard drives 191
protecting data
alternative methods 182
RAID methods 175

R

rack stability, warning xviii
RAID 0 (no fault tolerance) 176
RAID 1 (drive mirroring) 177
RAID 1+0 (data striping and drive mirroring) 177
RAID 5 (distributed data guarding) 179
RAID ADG (advanced data guarding) 180
RAID levels, changing for a LUN with the CLI 150
RAID methods
comparison with alternative fault-tolerance methods 182
software-based 182
summary of features 180
RAID methods, See also fault tolerance
rear view features 6
rear view, MSA1000, illustrated 6
rebuild
automatic data recovery 188
time 188
factors affecting 188
recognizing hard drive failure 185
reconfiguring a controller 101
recovery ROM 229
regulatory compliance identification numbers 163
regulatory compliance notices 164
related documentation xiv
removing
a spare drive 110
hard drive 189
replacing
batteries 67
failed drive 189
hard drive 189
resources, automatic data recovery 188
restrictions
moving arrays 191
moving drives 191
ROM
cloning 9
recovery 9

S

SCSI I/O Module with an Integrated Environmental Monitoring Unit (EMU)
described 15
functions of 15
illustrated 15
indicators, illustrated 56
SCSI ID assignments
4314 storage enclosure 233
4354 storage enclosure 234
MSA1000 drive shelf 232
Selective Storage Presentation (SSP), with the CLI 152, 156
serial connection, setting up 121
setup, overview of 34
setup, steps of 34
SFP
inserting the replacement transceiver 78
software components supported
Linux 26
NetWare 28
OpenVMS 29
Tru64 UNIX 31
Windows 25

software components, summary of supported 24
spare drives
  adding with the CLI 145
  defined 175
  deleting with the CLI 147
specifications 169
SSP
  defined 14
  diagram 14
standby power 8
status indicators 185
stripe size
  migrating 109
striping data, defined 173
switches
  power 8, 51
symbols
  in text xvi
  on equipment xvi

T

technical support, hp xviii
text symbols xvi
transceiver
  inserting the replacement SFP 78
troubleshooting
  cabling 187
  hard drive problems 185
troubleshooting See also Controller display messages
Tru64 UNIX environments
  specific configuration needs 31
  supported HBA 31
  supported operating systems 31
  supported software components 31
Tru64 UNIX environments, information about 31

U

unrecoverable disk error message 186

V

view
  front of the MSA1000 5
  rear of the MSA1000 6

W

warning
  rack stability xviii
websites
  hp storage xix
Windows environments
  specific configuration needs 26
  supported HBA 26
  supported operating systems 25
  supported software components 25
Windows environments, information about 25